Internet Draft
Expiration: August 8, 2004
Anti-Spam Research Group

J. Levine, A. DeKok, et al.
Taughannock Networks and IDT, Inc.

February 8, 2004

Lightweight MTA Authentication Protocol (LMAP) Discussion and Comparison

draft-irtf-asrg-lmap-discussion-01.txt

Status of this Memo

Copyright Notice

Abstract

Lightweight MTA Authentication Protocol (LMAP) is the general term for a a family of proposed protocols to help address the spam problem by better authenticating mail senders. This document discusses the applicability, and the costs and benefits of wide-spread deployment of the protocol, and compares the various LMAP proposals.

Table of Contents

1. Introduction

The Anti Spam Research Group (ASRG) was chartered to address the spam problem. The LMAP family of protocols falls with the scope of "changes to existing applications" as described in the charter.

The intended audience for this document includes administrators and developers of DNS and SMTP systems. The audience are assumed to be familiar with the workings of SMTP [RFC2821] and DNS [RFC1034].

The document is organized as follows. Section 2 describes the problem statement and scope of the proposal. Section 3 deals with the most commonly raised concerns associated with LMAP, and prior variants thereof.

This document is intended to evolve, based on comments from the Anti-Spam Research Group (ASRG) mailing list. It is certain that the current draft is incomplete and entirely possible that it is inaccurate. Hence, comments are eagerly sought, preferably in the form of suggested text changes, and preferably on the ASRG mailing list, at <asrg@ietf.org>.

These protocols are experimental and is not suitable for widespread deployment or production use.

We wish to remind readers that implementation and use of LMAP is entirely optional and not required to operate SMTP services. This document contains minor updates to the semantics of parts of RFC 2821.

Readers are further reminded that recipients have the right to refuse any communication from anyone, for any reason.

1.1. Summary of the Protocols

LMAP is based on two concepts: publication of authentication data by a domain, and application of that data by a recipient MTA. The combination of these concepts permits SMTP recipients to establish more reliably whether mail putatively from a domain is actually from that domain and that there is a responsible contact in case of questions or problems with the domain's mail.

The data published by a domain includes statements as to which IP's are permitted to originate mail from the domain in SMTP EHLO/HELO and MAIL FROM.

SMTP recipients can look up the authentication data when a domain name is used in EHLO/HELO and/or MAIL FROM. The recipient can then choose to apply the requested policy to the message. The result is that messages are delivered only when all parties consent to their delivery. The message originator must have the consent of the domain to claim association with that domain, and the recipient can verify that that consent exists. After all, if the domain does not consent to an originator claiming association with it, there are few reasons why a recipient would choose to accept that non-consensual message.

The method of establishing whether an IP address can send messages for a domain is similar to the usage of Dialup User Lists (DUL) such as the MAPS DUL. With those methods, a network provider publishes a list of IP addresses which have been assigned to dial-up users. SMTP recipients may query

such lists, and assume that the existence of an IP address on the list means that the network provider did not intend SMTP traffic to originate from that IP.

Similarly, methods such as sender callback attempt to discover the implied intentions of a domain, by performing certain queries to that domain, but such methods only perform a weak substitute for authentication.

This proposal makes all such authentication explicit, through the use of published data by a domain. Publication makes the Internet more open. Less information is hidden, and fewer erroneous implications are arrived at.

Note that recipients should also publish that they use/enforce LMAP. Receiving mail transport agents using protocols with the ability to advertise capabilities should advertise a capability to the sender that informs the that the sender that the receiver will check the incoming IP address with LMAP. It is to the advantage of all parties for a sender that will not be able to pass LMAP authentication to be able to discover the fact as early as possible and abort the transmission.

This verification scheme is weaker than cryptographic systems but stronger than the current SMTP model.

These proposals change the semantics of the MAIL FROM command as defined in RFC 2821, section 3.3. to imply that the domain in the source mailbox is also the responsible party for sending the message, and thus must be verified.

1.2. Interpretation of LMAP Data

Recipient MTAs are free to interpret LMAP data as they wish. Possibilities range from rejecting email with a 550 error code to using LMAP data as one input to a multi-criterion filter. Domains may also optionally use LMAP data to whitelist or give higher passing values for email in their filters.

E-mail from LMAP domains that do not publish LMAP data should NOT be rejected since there is no requirement that domains do so, and many will not, either for policy reasons or from lack of resources. E-email from non-LMAP domains should be treated as e-mail is treated today.

The local policy decisions remain with the recipient's MTA. Readers are reminded that recipients have the right to refuse any communication from anyone, for any reason.

2. Problem Statement and Scope

SMTP, as defined in RFC 821 and RFC 2821, provides no authentication at all when transferring mail. In the early years of the Internet, when the net was small and all networks

enforced acceptable use policies, that didn't present a problem. Since then, the size of the net has grown enormously, the number of networks has also grown, and policies vary greatly. As a result, it has become possible and increasingly common for irresponsible users to forge addresses without the permission of the legitimate owners of those addresses. The vast majority of mail with forged addresses is abusive for reasons beyond the forged addresses. Hence, measures to limit address forgery are likely also to limit the abuses associated with forgery.

LMAP attacks the forgery problem by checking that the host from which the message was sent is authorized to send mail using the a domain in the message's envelope. While this only deters a single category of forgery, the category it attacks is a large one, and by ensuring that at least one address is valid, recipients have both a reliable channel back to the domain's management, and provides a useful criterion for mail filtering, as described in the previous section.

2.1. Types of forgery

Mail forgery is associated with several varieties of abusive e-mail.

* Senders of junk email (the largest category of spam), often forges return addresses. It does so to make it harder to determine the responsible party, making it harder to tell to whom to complain. It also does so to evade filters, either by pretending to be a sender on a recipient's whitelist, or to pretend not to be a sender on a recipient's blacklist.

* In account fraud, also known as ``phishing'', a sender poses as a person or organization with whom the recipient has a business relationship, or would like to have a business relationship. The sender does so in order to trick the recipient into revealing personal information that the sender can use for identity theft or related crimes.

* In a ``joe job,'' a sender sends out abusive mail and forges the address of an unrelated party. The goal is to discredit the party whose address is forged.

* Viruses, trojans, worms, and related automated malware use forged return addresses to trick recipients into accepting or opening messages with hostile active content. The goal is to get the recipients to run the active content, thereby propagating the malware.

2.2. Choice of data to authenticate

When a message is sent via SMTP, the recipient MTA has a variety of items that it could use to authenticate the mail sender. In the order that they are available to the recipient

MTA they are:

2.2.1 IP address

The MTA can easily determine the IP address from which a mail message is sent. Many parties publish DNS blacklists and DNS whitelists (DNSBLs and DNSWLs) with lists of IP addresses from which the publishers recommend that recipient MTAs reject or accept incoming mail. DNSBLs are reasonably effective at identifying IP address ranges of networks with a history of sending abusive mail, of IP addresses of specific hosts with a history of sending abusive mail, typically because they are controlled by a trojan or other hostile software, and ranges of IP addresses assigned hosts whose users generally send mail via their ISP's mail servers and are not expected to send mail directly.

The MTA Mark[N3] and Selective Sender[N5] proposals are an extremely simple form of LMAP in which the owner of a range of IP addresses (or more precisely, the party that controls a range's reverse DNS, which should be the same as the network owner but is not always in practice) can identify which hosts in the network should be sending e-mail and which shouldn't. MTA Mark provides a further facility to publish a contact e-mail address or URL that recipients can use in case of questions or complaints about mail from a specific address.

2.2.2 HELO/EHLO name

Section 4.1.1.1 of RFC 2821 defines the HELO and EHLO commands which SMTP clients should use at the beginning of each SMTP session. The argument to HELO or EHLO is a domain or a domain literal that should identify the sending host.

2.2.3 Return path

Section 4.1.1.2 of RFC 2821 defines the MAIL command which SMTP clients use to provide the reverse path for a message. The reverse path is usually an e-mail address, but in the case of a bounce message (or a message pretending to be a bounce message), the reverse path can be empty.

2.2.4 Message header fields

RFC 2822 defines a list of message header fields including From:, Sender:, Resent-From:. and Resent-Sender: that include the e-mail addresses of the party or parties responsible for a particular message. The SMTP server receives the message headers and message body in a block after accepting the DATA

command defined by section 4.1.1.4 of RFC 2821. SMTP makes no provision for sending just the headers or a partial message; if an SMTP server accepts any of the headers at all, it must

be prepared to accept the entire message.

2.2.5 Comparison

All of these items, IP address, HELO/EHLO argument, return path, and message headers, can be used for various kinds of authentication.

The IP address is the cheapest to use, since it can be determined even before the SMTP session begins, but reveals the least information about the mail sender. There is no connection between an IP address and a responsible domain or e-mail address, unless there is MTA Mark contact information for the IP. Even if an IP address is a legitimate mail sending host, the IP address provides no basis to distinguish between mail sent by the host's legitimate users and mail sent by hostile users or malware with access to that host. The IP address gives no information about what domains have authorized mail to be sent from that host.

The HELO/EHLO string, if it is a domain name, can be validated using an LMAP-like scheme such as DRIP [N2] to check that the host at the connecting IP is authorized to use the domain name it is offering. Like IP validation, HELO/EHLO validation doesn't provide a contact e-mail address. The HELO/EHLO address should be the host's name, such as MAIL.EXAMPLE.COM, even if the mail it sends all has addresses in the domain EXAMPLE.COM. Although the address Postmaster@EXAMPLE.COM would exist, the address Postmaster@MAIL.EXAMPLE.COM may well not exist if no legitimate mail is sent with MAIL.EXAMPLE.COM return addresses. If the HELO/EHLO string is a domain literal, it can't be verified other than by IP address verification, and it's unlikely to be useful for contact e-mail since few hosts accept mail to domain literal addresses. RFC 2821 permits the SMTP server to validate a domain address provided to EHLO or HELO by doing a DNS lookup to see if the IP address matches, but does not permit the server to reject mail if the validation fails. HELO/EHLO-based LMAP would modify RFC2821 by allowing the server to reject mail based on HELO/EHLO validation failure.

The return path, if its domain is verified by LMAP, is likely to be a valid contact address for the message. Since LMAP only looks at the domain, not the full address, it's still possible that the mailbox isn't valid, in which case the contact address might be abuse@(domain) or Postmaster@(domain). Return paths can legitimately be empty, in which case the server would either have to skip LMAP

validation, or use HELO/EHLO or message header data instead.

A message header address, if verified by LMAP, is also likely to be a valid contact address for the message. Any valid message has at least one address in a From: or Sender: line.

Although any or all of these items could be used for message validation, the LMAP proposals use the return path as a compromise between the best quality data and efficiency of operation. A verified return path, unlike IP or HELO/EHLO data, is very likely to provide a specific responsible address and responsible domain for a message. Existing SMTP implementations tend to preserve the return path throughout the delivery process so it can be used at any stage that a contact address is needed. The HELO/EHLO domain, on the other hand, is generally preserved only in a Received: header, and although RFC 2821 specifies the way that the domain is stored in that header, MTAs do not all conform to it.

A verified message header address is also high-quality data, but the cost of extracting the header addresses is much higher than for using return path or EHLO/HELO domains because of the cost of receiving the entire message before deciding whether it passes LMAP validation and of parsing the headers, which is otherwise not needed at SMTP time. If an SMTP server is going to receive the entire message anyway, it might be appropriate to apply other more powerful cryptographic signature verification instead of or in addition to LMAP.

2.3. DNS based attacks

All versions of LMAP use the DNS to distribute the data against which mail is authenticated. This makes the DNS the critical resource required by all of these proposals. Insecurities in the DNS, as described in [7], could allow hostile parties to page forged authentication information into the DNS. Packet floods and other denial of service attacks against DNS servers could make it impossible for LMAP clients to obtain LMAP authentication data.

3. Common Concerns with LMAP

This section describes the most common concerns raised about LMAP, and responds in detail to those concerns.

3.1. LMAP and the end to end nature of the Internet

Concerns have been raised that this proposal negatively affects the "end to end" nature of the Internet. LMAP does not change SMTP, except for changing the semantics of the mailbox used in MAIL FROM command. The end to end nature of SMTP is therefore unchanged. What this proposal does offer is a way to hold the originating end of an SMTP session accountable for any association it alleges it has with a

domain. Claims that this accountability is an unwarranted restriction on the "end to end" nature of the Internet should consider:

1) SMTP originators who wish to be unaccountable for their

behavior are called "spammers". The intention of this proposal is to address the problem of spam, not to condone it.

2) SMTP originators who wish to force their messages onto recipients, despite the recipients desire not to receive them, are also called "spammers". If a recipient chooses to request that a sender be publicly accountable for his behavior and the sender refuses, then the recipient is free to reject or discard any messages from the sender.

The lack of accountability is a major technical reason why spam is such a problem.

This proposal extends the end-to-end principles on which the Internet was built, because it allows each end to publish its policy, and to discover the others policy. Sharing of information enhances trust, and permits the discovery of problems related to Man in the Middle (MITM) issues.

3.2. Roaming Users and LMAP

Another concern raised about LMAP is that it will negatively affect roaming users, that is, users not connected to their usual or home network. The main concern of roaming users is that the deployment of LMAP will break the "end to end" nature of the Internet.

The response to those concerns can best be summarized as follows:

Stopping mail forgery requires every one of them to give up forging.

The practices of roaming users currently require that the SMTP recipient do significant amounts of work to authenticate or filter their messages. Recipients in return request that the roaming user (and the alleged originating domain) share some of that work. This response serves as the foundation for the design of LMAP.

If the roaming user is unwilling to share the work of demonstrating accountability, then the recipient is as always free to reject any communication with that roaming user.

Existing practice includes a variety of methods through which roaming users may send email messages in circumstances where LMAP is widely deployed.

3.2.1 SUBMIT (port 587)

That is, SMTP on another port [RFC 2476]. Roaming users can use SUBMIT to send messages to their home provider, which then sends those messages to the final recipients via SMTP. This

method has all of the benefits of SMTP, with none of the drawbacks of recipients being responsible for authenticating roaming end-users.

The primary cost or delay associated with this method is deployment in mail client software.

3.2.2 SMTP relaying to the home provider

Since many network providers currently block outgoing SMTP traffic (on port 25), this option is not universally available. The roaming users are then in the awkward circumstance of having their attempt to behave like spammers blocked, by an attempt to prevent spam.

3.2.3 Publish LMAP information

Roaming users may update LMAP information for their domain through Dynamic DNS (DDNS). Any messages they send will then pass LMAP criteria, subject to DNS propagation delays.

Roaming users can also publish a policy though LMAP that any IP address on the Internet is permitted to claim association with their domain. Administrators who publish such information for their domain should be aware that this practice will open them up to spammers claiming association with their domain. For this reason, we do not recommend such practices.

3.2.4 Virtual Private Networks (VPNs)

With this solution, roaming users allow their home provider to authenticate them, and any SMTP traffic is sent through a secure tunnel. That traffic then appears to issue from the network of the home provider, where LMAP information may easily be published and maintained.

3.2.5 Other message delivery systems

Bi-directional POP, IMAP, Webmail, etc. all exist, and are sub-optimal. But they work.

They also have the added benefit that they are not required to scale with the Internet. Rather, they scale with the number of users at a domain. So it is not the problem of the rest of the Internet to deal with those issues, but instead the domain with roaming users.

3.3. Message relaying and forwarding are affected by LMAP

Mail forwarders have traditionally left the sender envelope untouched. "Forwarding" is used in the sense of Unix user .forward and forwarding services such as those provided by pobox.com and ieee.org.

Let us examine a situation where an LMAP conformant domain A sends a message to address B which forwards the message to LMAP conformant recipient C using the original sender address from A. If the B->C forward had been set up without the consent of the recipient C, A's LMAP records would be checked by C's LMAP client, and the message would be correctly rejected.

If the recipient C did desire the B->C forwarding, possible changes to work with LMAP include:

1) B's MTA could rewrite the sender address to to one in B's domain.

2) the user B could alter the .forward to apply a return path in B's domain

3) the recipient C could provide a whitelist to C's MTA indicating that forwarded messages are expected to arrive for C from B.

LMAP conformant SMTP forwarders could implement a sender rewriting scheme or its equivalent. The technical details of doing so appear simple in most popular mail systems.

3.4. Kiosk, greeting card, and mail-a-link systems

Many web sites offer a facility to mail content from the web site to a third party, with the web user's return address. The content may be a greeting card, a magazine article, or a message entered by the user. Few of these sites do any validation of the sender's address, although they tend to be rate limited or inherently slow enough that they're not useful for sending out spam, but since users can enter any return address, the mail they send is technically indistinguishable from mail with forged return addresses.

The most straightforward way to make such systems comply with LMAP would be for them to use their own domain in the return address, while using the user's entered address on the From: line.

4. Comparison of Proposals

Several different varieties of LMAP have been proposed in recent months. They include:

RMX (Hadmut Danisch) [N5]

DMP (Designated Mailers Protocol, Gordon Fecyk) [N1]

SPF (Sender Permitted From, Mark Lentczner and Meng Weng Wong) [N7]

FSV (Flexible Sender Validation, John Levine) [N3]

Two simpler but similar proposals are:

MM (MTA Mark, Markus Stumpf and Steff Hoehne) [N4]

SS (Selective Sender, John Levine) [N6]

They both identify IP addresses as mail senders or not, without asserting anything about what e-mail addresses should originate mail from what address.

Each of these proposals is a work in progress, hence the reader must refer to the latest defining document for each to learn the exact details of each. These comparisons are not intended to recommend one proposal over another, but rather to highlight the differences among them.

4.1. DNS record types

RMX: defines two new record types, RMX and APL

SPF: uses standard TXT records

DMP: uses standard TXT records

FSV: uses standard TXT and A records

MM: uses standard TXT records

SS: uses standard TXT records

4.2. Record structure

RMX: RMX records have multiple subtypes, APL records are a list of CIDR ranges

SPF: TXT records contain syntax that require parsing

DMP: TXT records contain fixed strings

FSV: TXT records contain ASCII CIDR ranges

MM and SS: TXT records contain fixed strings

4.3. Indirect addressing

RMX: RMX records can refer to RMX and APL records in other domains

SPF: syntax includes indirect references to SPF data in other domains

DMP: none other than CNAMES

FSV: none other than CNAMES

MM and SS: none other than CNAMES

4.4. Address index structure

(That is, how is the set of authorized IP addresses for a domain stored.)

RMX: Aggregated data for a domain are contained in APL records, or implicitly via host or MX records

SPF: Similar to RMX, with all data encoded in TXT records

DMP: Individual records per IP per domain

FSV: Both individual A records per IP per domain, and aggregate data in TXT record for a domain

DMP: Individual records per IP per domain

MM and SS: Individual records per IP address

4.5. Roaming and forwarding support

RMX: None

SPF: Domain's data can permit mail from anywhere in addition to listed IPs

DMP: HELO/EHLO name can be validated instead of envelope

FSV: same as DMP

MM and SS: Not an issue, doesn't attempt to validate domains

4.6. Number of queries

RMX: Data for domain retrieved in one or more queries, can be cached for future mail from the same domain

SPF: Similar to RMX with multiple levels of indirection possible

DMP: Single domain+IP query to validate a message. In case of failure, possible second query to see if domain announces DMP data.

FSV: Single query to fetch all of domain's data, or single domain+IP to validate a single message, with possible second

query to see if domain announces FSV data.

MM and SS: One per IP address

4.7. Extensibility

RMX: not addressed

SPF: version number in data and extensible record syntax allow for additional definitions

DMP: not addressed, new fixed strings could be defined

FSV: not addressed, new fixed strings could be defined

MM and SS: not addressed, new fixed strings could be defined

5. Privacy Considerations

This proposal does not examine message contents, or user identities in MAIL FROM. It therefore has no privacy considerations which affect those fields.

The largest effects of this proposal on privacy are in three areas: end users, publication of network infrastructure, and traffic analysis.

5.1. End Users

This proposal affects the privacy of end users in two ways. First, it permits recipients to associate user identities or message contents with a domain that is accountable for the message. Second, it prevents users from sending mail pseudonymously, by using an address in another domain, or in a non-existent domain.

The first effect on privacy has little impact. The user sending the message is already claiming association with a domain, so there is no loss of privacy if that association is verified by the message recipient. Also, as noted previously, this proposal does not prevent users from fraudulently claiming to be another user within a domain.

The second effect on privacy may be considered undesirable by some observers. True anonymity through the practice of forged association with domains, forged email addresses, and by sending email through hijacked or trojaned systems will become more difficult. This sort of anonymity is highly correlated with spam, however, and is precisely the kind of abuse that this proposal attempts to prevent.

Users who wish anonymity may gain it through accountable mechanisms. Throw-away accounts at reputable network providers may be created and paid for in cash under an assumed name, for example. Other organizations will guarantee a degree of anonymity (more realistically called shelter from

others), if certain requirements are met.

We suggest that accountable methods of creating anonymity be used, rather than unaccountable methods. One individual's desire for anonymity does not, and should not, require the rest of the Internet to accept large volumes of spam.

5.2. Network Infrastructure

Publication of LMAP information results in a readily available list of IP addresses of hosts authorized to send messages associated with a domain. These lists yield information about the network structure, and business relationships, and presents hostile parties with a list of targets to attempt to compromise.

However, such information is often already publicly accessible through other means. Anyone communicating with individuals at a domain may readily obtain this information, and share it with anyone else. Business relationships have been discovered, for example, prior to "official" public announcement, by examining DNS records. Nearly all such "private" information about network structure and relationships may therefore be described as already being readily available.

If such information is to be kept secret, it is the users responsibility to send messages in such a way as to keep that information private.

5.3. Traffic analysis

Any LMAP aware MTA and DNS server requires additional network traffic beyond that required by SMTP. This traffic may be analyzed in order to verify that two parties are communicating, or that a particular message has been received. The additional traffic may still be analyzed in this manner, even if the SMTP session is encrypted.

However, many MTAs already query MX and A records of a domain after receiving a MAIL FROM command, so the threat of this new traffic is minimal.

6. Security Considerations

This document describes common uses of LMAP, attacks on it, and defenses which may be implemented. While much of this document deals with security issues, it does not propose any standard, and therefore does not have any direct security

effects.

However, implementors and administrators of systems using LMAP

should be aware of the issues raised herein.

7. Informative References

[1] Bradner, S., "The Internet Standards Process -- Revision 3", BCP9, RFC 2026, October 1996.

[2] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[3] Crocker, D. "Technical Considerations for Spam Control Mechanisms", work in progress, http://brandenburg.com/specifications/draft-crocker-spam-techconsider-02.txt

[4] Hoffman, P. "SMTP Service Extension for Secure SMTP over TLS", RFC 2487, January, 1999

[5] Myers, J. "SMTP Service Extension for Authentication", RFC 2554, March, 1999.

[6] Klensin, J. (Ed) "Simple Mail Transfer Protocol", RFC 2821, April 2001.

[7] D. Atkins and R. Austein, "Threat Analysis Of The Domain Name System," Internet Draft internet-drafts/draft-ietf-dnsext-dns-threats-05.txt, Nov 2003.

8. Normative References

[N1] G. Feyck, "Designated Mailers Protocol (DMP)", Internet draft draft-fecyk-dmp-01.txt, December 2003.

[N2] R. S. Brand and L. Sherzer, "Designated Relays Inquiry Protocol (DRIP)," http://www.sherzer.net/draft-brand-drip-02.txt, October 2003,

[N3] J. Levine, "Flexible Sender Validation (FSV)", http://www.taugh.com/draft-levine-fsv-00.txt, February 2004.

[N4] M. Stumpf and S. Hoehne. "MTA Mark", http://www.space.net/~maex/draft-irtf-asrg-mtamark-00.txt, September 2003.

[N5] H. Danisch, "The RMX DNS RR and method for lightweight SMTP sender authorization ", IETF draft draft-danisch-dns-rr-smtp-03.txt, October 2003.

[N6] J. Levine, Selective Sender,

http://www.taugh.com/mp/ss.html, January 2004.

[N7] M. Lentczner and M. W. Wong, "Sender Authentication with Sender Permitted From (SPF) A Convention to Describe Hosts Authorized to Send SMTP Traffic", http://spf.pobox.com/draft-

mengwong-spf.02.9.6.txt, January 2004.

9. Authors' Addresses

John R. Levine
Taughannock Networks
PO Box 727
Trumansburg NY 14886 USA
E-mail: johnl@taugh.com
Phone: +1 607 330 5711

Alan DeKok
IDT Canada, Inc.
1575 Carling Ave.
Ottawa, ON K1G 0T3 Canada
Email: alan.dekok@idt.com
Phone: +1 613 724 6004 ext. 231

Full Copyright Statement

$Id: draft-irtf-asrg-lmap-discussion-00.n,v 1.5 2004/02/08 21:19:27 johnl Exp $