

Network Working Group  
Request for Comments: 4417  
Category: Informational

P. Resnick, Ed.  
IAB  
P. Saint-Andre, Ed.  
JSF  
February 2006

## Report of the 2004 IAB Messaging Workshop

### Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

### Copyright Notice

Copyright (C) The Internet Society (2006).

### Abstract

This document reports the outcome of a workshop held by the Internet Architecture Board (IAB) on the future of Internet messaging. The workshop was held on 6 and 7 October 2004 in Burlingame, CA, USA. The goal of the workshop was to examine the current state of different messaging technologies on the Internet (including, but not limited to, electronic mail, instant messaging, and voice messaging), to look at their commonalities and differences, and to find engineering, research, and architectural topics on which future work could be done. This report summarizes the discussions and conclusions of the workshop and of the IAB.

## Table of Contents

1. Introduction .....	3
2. Methodology .....	4
3. Issues .....	5
3.1. Authorization .....	5
3.2. Multiple Communication Channels .....	6
3.3. Negotiation .....	8
3.4. User Control .....	9
3.5. Message Transport .....	9
3.6. Identity Hints and Key Distribution .....	10
4. Recommendations .....	11
4.1. Authorization .....	11
4.2. Multiple Communication Channels .....	12
4.3. Negotiation .....	13
4.4. User Control .....	13
4.5. Message Transport .....	14
4.6. Identity Hints and Key Distribution .....	16
5. Security Considerations .....	16
6. Acknowledgements .....	16
Appendix A. Participants .....	17
Appendix B. Pre-Workshop Papers .....	18

## 1. Introduction

Current email infrastructure is a mixture of facilities to accomplish its task of end-to-end communications through a relay mesh. That mixture has come about as requirements have changed over the years. Discussions recur over the years, often including complaints that some desired features of email (such as internationalization, efficient encoding of structured data, trusted communication) are ill-served by the current infrastructure, or that some of the current infrastructure seems to be adversely affected by current problems on the Internet (most recently including problems such as spam, viruses, and lack of security infrastructure). For many years, the daunting task of revamping email infrastructure has been considered, with justifiably little enthusiasm for taking on such a task. However, there has been some recent informal discussion on the kinds of things that would be desirable in a "next generation" email.

At the same time, other messaging infrastructures (including those associated with "instant messaging" and "web logging") are currently being deployed that appear to address many of the above desired features and outstanding problems, while adding many features not currently considered part of traditional email (like prior-consent-based acceptance of messages). However, each of these technologies (at least in their current deployment) seem to lack some of the features commonly associated with email (such as selective and partial message delivery, queued multi-hop relaying, offline message management, and efficient non-textual content delivery).

The Internet Architecture Board (IAB) believed that the time was ripe to examine the current state of messaging technologies on the Internet and to see if there are areas of work that can be taken on to advance these technologies. Therefore, the IAB held a workshop on Internet messaging, taking some of the above issues as input, in order to formulate some direction for future study of the area of messaging.

The topic of messaging is broad, and the boundaries of what counts as messaging are not always well-defined. Rather than limit themselves to a philosophical discussion of the nature of messages, the workshop participants adopted the attitude of "we know it when we see it" and used as their primary examples such well-established types of messaging as email and instant messaging (IM), while also discussing more "peripheral" types of messaging such as voice messaging and event notifications. (Message queuing systems with guaranteed delivery and transactional integrity, such as those used in enterprise workflow engines and some "web services" architectures, were operationally if not intentionally out of scope.) The participants worked to discover common themes that apply to all the

types of messaging under consideration. Among the themes identified were the following:

- o Authorization of senders and recipients
- o Negotiation of messaging parameters
- o Consent models and privacy
- o Identity hints, reputation, and key distribution
- o Cross-protocol unification of messaging models
- o Enabling greater user control over messaging
- o Transport issues (unreliable links, push/pull, etc.)
- o Message organization (e.g., conversations and threading)

Purposely missing from the foregoing list is the topic of unsolicited commercial email or unsolicited bulk email (UCE or UBE, colloquially known as "spam") and analogous communications in other messaging environments such as instant messaging ("spim") and Internet telephony ("spit"). While this topic was an impetus for the IAB's holding the workshop, it was kept off the workshop agenda due to concerns that it would crowd out discussion of other messaging-related issues. The more general topics of authorization and identity were thought to be broad enough to cover the architectural issues involved with spam without devolving into more unproductive discussions.

This document is structured so as to provide an overview of the discussion flow as well as proposed recommendations of the workshop. Section 3 summarizes the discussions that occurred during the workshop on various topics or themes, while Section 4 provides an overview of recommended research topics and protocol definition efforts that resulted from the workshop. Section 5 provides some perspective on the security-related aspects of the topics discussed during the workshop. Appendix B lists the pre-workshop topic papers submitted by workshop participants as background for the workshop discussions.

## 2. Methodology

Prior to the workshop, brief topic papers were submitted to set the context for the discussions to follow; a list of the papers and their authors is provided in Appendix B of this document.

During the workshop itself, discussion centered on several topics or themes, as summarized in the following sections. Naturally, it was not possible in a two-day workshop to treat these topics in depth; however, rough consensus was reached on the importance of these topics, if not always on the details of potential research programs and protocol standardization efforts that might address the issues

raised. It is hoped that these summaries will inspire work by additional investigators.

The in-workshop discussions quite naturally fell into three kinds of "tracks": (1) possible engineering tasks to recommend to the IESG and other standardization groups, (2) "blue sky" research topics to recommend to the IRTF and other researchers, and (3) general architectural or "framework" issues for consideration by both engineers and researchers alike. After a full-group discussion each morning to identify possible topics for more in-depth investigation, participants self-selected for involvement in one of three "break-out" sessions. Toward the end of each day, the full groups reconvened, gathered reports from the break-out discussion leaders, and attempted to come to consensus regarding lessons learned and recommendations for further research. The results of the two-day workshop therefore consist of discussion issues and research/engineering recommendations related to the six topics described in this report.

### 3. Issues

#### 3.1. Authorization

It is one thing for a sender to send a message, and another thing for the intended recipient to accept it. The factors that lead a recipient to accept a message include the identity of the sender, previous experience with the sender, the existence of an ongoing conversation between the parties, meta-data about the message (e.g., its subject or size), the message medium (e.g., email vs. IM), and temporal or psychological factors. Authorization or acceptance applies most commonly at the level of the message or the level of the sender, and occasionally also at other levels (conversation thread, medium, sender domain).

Traditionally, sender authorization has been handled by recipient-defined block and allow lists (also called "blacklists" and "whitelists"). Block lists are of limited value, given the ease of gaining or creating new messaging identities (e.g., an email address or IM address). Allow lists are much more effective (since the list of people you like or want to communicate with is smaller than the large universe of people you don't), but they make it difficult for a sender to initiate communication with a new or previously unknown recipient. The workshop participants discussed several ways around this problem, including reputation systems and better ways for one person to introduce another person to a third party (e.g., through signed invitations).

Reputation systems may be especially worthy of future research, since they emulate a pattern that is familiar from real life. (It may also be valuable to distinguish between (1) reputation as the reactive assessment of a sender created by one or more recipients based on message history and (2) accreditation as a proactive assessment provided by trusted third parties.) Reputation might be based on summing an individual's "scores" provided by recipients on the network. (Naturally, the more important reputation becomes, the more bad actors might attempt to sabotage any given reputation system, so that a distributed as opposed to centralized system might be more desirable.) The actions taken by any given recipient based on the sender's reputation would not necessarily be limited to a simple allow/deny decision; more subtle actions might include placing messages from individuals with lower reputation scores into separate inboxes or redirecting them to other media (e.g., from IM to email).

### 3.2. Multiple Communication Channels

It is a fact of life that many people use multiple forms of messaging channels: phone, email, IM, pager, and so on. Unfortunately, this can make it difficult for a sender or initiator to know the best way to contact a recipient at any given time. One model is for the initiator to guess, for example, by first sending an email message and then escalating to pager or telephone if necessary; this may result in delivery of redundant messages to the recipient. A second model is for the recipient to publish updated contact information on a regular basis, perhaps as one aspect of his or her presence; this might enable the initiator to determine beforehand which contact medium is most appropriate. A third model is for the recipient to use some kind of "unifier" service that enables intelligent routing of messages or notifications to the recipient based on a set of delivery rules (e.g., "notify me via pager if I receive a voicemail message from my boss after 17:00").

The workshop participants did not think it necessary to choose between these models, but did identify several issues that are relevant in unifying or at least coordinating communication across multiple messaging channels:

- o While suppression of duplicate messages could be enabled by setting something like a "seen" flag on copies received via different messaging media, in general the correlation of multi-channel, multi-message exchanges is not well supported by existing standards.
- o A recipient could communicate his or her best contact mechanism to the initiator by explicitly granting permission to the initiator, perhaps by means of a kind of "authorization token".

- o It may be worthwhile to define frameworks or protocols for recipient-defined delivery rules. Currently, routing decisions tend to be made mostly by the sender through the choice of a messaging channel, but in the future the recipient may play a larger role in such decisions.
- o The logic behind contact publication needs to be explored, for example, whether it is an aspect of or extension to presence and whether contact addresses for one medium are best obtained by communicating in a different medium ("email me to get my mobile number").

A multiplicity of delivery channels also makes it more complex for a senders to establish a "reliable" relationship with a recipient. From the sender's point of view, it is not obvious that a recipient on one channel is the same recipient on another channel. How these recipient "identities" are tied together is an open question.

Another area for investigation is that of recipient capabilities. When the sender does not have capability information, the most common result is downgrading to a lowest common denominator of communication, which seriously underutilizes the capabilities of the entire system. Previous standards efforts (e.g., LDAP, Rescap, vCard, Conneg) have attempted to address parts of the capability puzzle, but without great success.

The existing deployment model uses several out-of-band mechanisms for establishing communications in the absence of programmatic capabilities information. Many of these mechanisms are based on direct human interaction and social policies, which in many cases are quite efficient and more appropriate than any protocol-based means. However, a programmatic means for establishing communications between "arms length" parties (e.g., business-to-business and business-to-customer relationships) might be very beneficial.

Any discussion of relationships inevitably leads to a discussion of trust (e.g., "from what kinds of entities do I want to receive messages?"). While this is a large topic, the group did discuss several ideas that might make it easier to broker communications within different relationships, including:

- o Whitelisting is the explicit definition of a relationship from the recipient's point of view, consisting of a list of senders with whom a recipient is willing to engage in conversation. While allow lists can be a workable solution, they are a relatively static authorization scheme.

- o Token-based authorization enables the recipient to define a one-time or limited-time relationship with a sender. The issuer possesses a token that grants a limited-time right to communicate with the recipient. This is a more dynamic authorization scheme.
- o Rule-based authorization involves an algorithmic assessment of the viability of a relationship based on a wide set of criteria. This is a more general authorization scheme that can incorporate both allow lists and tokens, plus additional evaluation criteria such as message characterization and issuer characterization.

### 3.3. Negotiation

In the area of negotiation, the workshop participants focused mainly on the process by which a set of participants agree on the media and parameters by which they will communicate. (One example of the end result of such a "rendezvous" negotiation is a group of colleagues who agree to hold a voice conference, with a textual "groupchat" as a secondary communications channel.) In order to enable cross-media negotiation, it may be necessary to establish a bridge between various identities. For example, the negotiation may occur via email, but the communication may occur via phone, and in order to authorize participants the conference software needs to know their phone numbers, not their email addresses. Furthermore, the parameters to be negotiated may include a wide variety of aspects, including:

- o Prerequisites for the communication (e.g., distribution of a "backgrounder" document).
- o Who will initiate the communication.
- o Who will participate in the communication.
- o The primary "venue" (e.g., a telephone number that all participants will call).
- o One or more secondary venues (e.g., a chatroom address).
- o Backup plans if the primary or secondary venue is not available.
- o The topic or topics for the discussion.
- o The identities of administrators or moderators.
- o Whether or not the discussion will be logged or recorded.
- o Scheduling of the event, including recurrence (e.g., different instances may have different venues or other details).

Indeed, in some contexts it might even be desirable to negotiate or re-negotiate parameters after communication has already begun (e.g., to invite new participants or change key parameters such as logging). While the workshop participants recognized that in-depth negotiation of a full set of parameters is likely to be unnecessary in many classes of communication, parts of a generalized framework or protocol for the negotiation of multiparty communication might prove useful in a wide range of applications and contexts.



### 3.4. User Control

A common perception among "power users" (and, increasingly, average users) on the Internet is that messaging is not sufficiently under their control. This is not merely a matter of unsolicited communications, but also of managing multiple messaging media and handling the sheer volume of messages received from familiar and unfamiliar senders alike. Currently, individuals attempt to cope using various personal techniques and ad hoc software tools, but there may be an opportunity to provide more programmatic support within Internet protocols and technologies.

One area of investigation is message filtering. Based on certain information -- the identity of the sender and/or recipient(s), the sender's reputation, the message thread or conversational context, message headers, message content (e.g., the presence of attachments), and environmental factors such as time of day or personal mood -- a user or agent may decide to take one of a wide variety of actions with regard to a message (bounce, ignore, forward, file, replicate, archive, accept, notify, etc.). While it is an open question how much formalization would be necessary or even helpful in this process, the workgroup participants identified several areas of possible investigation:

- o Cross-media threads and conversations -- it may be helpful to determine ways to tag messages as belonging to a particular thread or conversation across media (e.g., a forum discussion that migrates to email or IM), either during or after a message exchange.
- o Communication hierarchies -- while much of the focus is on messages, often a message does not stand alone but exists in the context of higher-level constructs such as a thread (i.e., a coherent or ordered set of messages within a medium), a conversation (i.e., a set of threads that may cross media), or an activity (a set of conversations and related resources, such as documents).
- o Control protocols -- the workgroup participants left as an open question whether there may be a need for a cross-service control protocol for use in managing communications across messaging media.

### 3.5. Message Transport

Different messaging media use different underlying transports. For instance, some messaging systems are more tolerant of slow links or lossy links, while others may depend on less loss-tolerant transport mechanisms. Integrating media that have different transport profiles can be difficult. For one, assuming that the same addressing

endpoint represents the same entity over time may not be warranted (it is possible that further work in identifying, addressing, and discovering endpoints may be appropriate, even at the URI level). It is also possible that the same endpoint or entity could be available via different transport mechanisms at different times, or even available via multiple transports at the same time. The process of choosing an appropriate transport mechanism when there are multiple paths introduces addressing issues that have not yet been dealt with in Internet protocol development (possible heuristics might include predictive routing, opportunistic routing, and scheduled routing). For links that can be unreliable, there may be value in being able to gracefully restart the link after any given failure, possibly by switching to a different transport mechanism.

Another issue that arises in cross-media and cross-transport integration is synchronization of references. This applies to particular messages but might also apply to message fragments. It may be desirable for some message fragments, such as large ancillary data, to be transported separately from others, for example small essential text data. Message fragments might also be forwarded, replicated, archived, etc., separately from other parts of a message. One factor relevant to synchronization across transports is that some messaging media are push-oriented (e.g., IM) whereas others are generally pull-oriented (e.g., email); when content is pushed to a recipient in one medium before it has been pulled by the recipient in another medium, it is possible for content references to get out of sync.

If message fragments can be transported over different media, possibly arriving at separate times or through separate paths, the issue of package security becomes a serious one. Traditionally, messages are secured by encrypting the entire package at the head end and then decrypting it on the receiving end. However, if we want to allow transports to fragment messages based upon the media types of the parts, that approach will not be feasible.

### 3.6. Identity Hints and Key Distribution

While it is widely recognized that both message encryption and authentication of conversation partners are highly desirable, the consensus of the workshop participants was that current business and implementation models in part discourage deployment of existing solutions. For example, it is often hard to get new root certificates installed in clients, certificates are (or are perceived to be) difficult or expensive to obtain, one-click or zero-click service enrollment is a worthy but seemingly unreachable goal, and

once one has created a public/private key pair and certified the public key, it is less than obvious how to distribute that certificate or discover other people's certificates.

One factor that may make widespread message encryption more feasible is that email, instant messaging, and Internet telephony have quite similar trust models. Yet the definition of communication differs quite a bit between these technologies: in email "the message is the thing", and it is a discrete object in its own right; in telephony the focus is on the real-time flow of a conversation or session rather than discrete messages; and IM seems to hold a mediate position since it is centered on the rapid, back-and-forth exchange of text messages (which can be seen as messaging sessions).

Another complicating factor is the wide range of contexts in which messaging technologies are used: everything from casual conversations in public chatrooms and social networking applications, through communications between businesses and customers, to mission-critical business-to-business applications such as supply chain management. Different audiences may have different needs with regard to messaging security and identity verification, resulting in varying demand for services such as trusted third parties and webs of trust.

In the context of communication technologies, identity hints -- shared knowledge, conversational styles, voice tone, messaging patterns, vocabulary, and the like -- can often provide more useful information than key fingerprints, digital signatures, and other electronic artifacts, which are distant from the experience of most end users. To date, the checking of such identity hints is intuitive rather than programmatic.

#### 4. Recommendations

##### 4.1. Authorization

The one clearly desired engineering project that came out of the authorization discussion was a distributed reputation service. It was agreed that whatever else needed to be done in regard to authorization of messages, at some point the recipient of the message would want to be able to check the reputation of the sender of the message. This is especially useful in the case of senders with whom the recipient has no prior experience; i.e., using a reputation service as a way to get an "introduction to a stranger". There was clearly a need for this reputation service to be decentralized; though a single centralized reputation service can be useful in some contexts, it does not scale to an Internet-wide service.

Two potential research topics in authorization were discussed. First, a good deal of discussion centered around the use of whitelists and blacklists in authorization decision, but it was thought that research was necessary to examine the relative usefulness of each of the approaches fully. It was clear to the participants that blacklists can weed out known non-authorized senders, but do not stop "aggressive" unwanted senders because of the ease of simply obtaining a new identity. Whitelists can be useful for limiting messages to only those known to the recipient, but would require the use of some sort of introduction service in order to allow for messages from unknown parties. Participants also thought that there might be useful architectural work done in this area.

The other potential research area was in recipient responses to authorization decisions. Upon making an authorization decision, recipients have to do two things: First, obviously the recipient must dispatch the message in some way either to deliver it or to deny it. But that decision will also have side effects back into the next set of authorization decisions the recipient may make. The decision may feed back into the reputation system, either "lauding" or "censuring" the sender of the message.

#### 4.2. Multiple Communication Channels

Several interesting and potentially useful ideas were discussed during the session, which the participants worked to transform into research or engineering tasks, as appropriate.

In the area of contact information management, the workshop participants identified a possible engineering task to define a service that publishes contact information such as availability, capabilities, channel addresses (routing information), preferences, and policies. While aspects of this work have been attempted previously within the IETF (with varying degrees of success), there remain many potential benefits with regard to managing business-to-business and business-to-customer relationships.

The problem of suppressing redundant messages is becoming more important as the use of multiple messaging channels becomes the rule for most Internet users, and as users become accustomed to receiving notifications in one channel of communications received in another channel. Unfortunately, there are essentially no standards for cross-referencing and linking of messages across channels; standards work in this area may be appropriate.

Another possible engineering task is defining a standardized representation for the definition and application of recipient message processing rules. Such an effort would extend existing work on the Sieve language within the IETF to incorporate some of the concepts discussed above.

Discussion of token-based authorization focused on the concept of defining a means for establishing time-limited or usage-limited relationships for exchanging messages. The work would attempt to define the identity, generation, and use of tokens for authorization purposes. Most likely this is more of a research topic than an engineering topic.

Work on recipient rules processing and token-based authentication may be related at a higher level of abstraction (we can call it "recipient authorization processing"). When combined with insights into authorization (see Sections 3.1 and 4.1), this may be an appropriate topic for further research.

#### 4.3. Negotiation

Discussion in the area of negotiation resulted mostly in research-oriented output. The session felt that participants in a conversation would require some sort of rendezvous mechanism during which the parameters of the conversation would be negotiated. To facilitate this, a "conversation identifier" would be needed so that participants could identify the conversation that they wished to participate in. In addition, there are at least five dimensions along which a conversation negotiation may occur:

- o The participants in the conversation
- o The topic for the conversation
- o The scheduling and priority parameters
- o The mechanism used for the conversation
- o The capabilities of the participants
- o The logistical details of the conversation

Research into how to communicate these different parameters may prove useful, as may research into the relationship between the concepts of negotiation, rendezvous, and conversation.

#### 4.4. User Control

A clear architectural topic to come out of the user control discussion was work on activities, conversations, and threads. In the course of the discussion, the user's ability to organize messages into threads became a focus. The participants got some start on defining threads as a semi-ordered set of messages, a conversation as

a set of threads, and an activity as a collection of conversations and related resources. The discussion expanded the traditional notion of a thread as an ordered tree of messages. Conversations can collect together threads and have them be cross-media. Messages can potentially belong to more than one thread. Threads themselves might have subthreads. All of these topics require an architectural overview to be brought into focus.

There is also engineering work that is already at a sufficient level of maturity to be undertaken on threads. Though there is certainly some simple threading work being done now with messaging, it is pretty much useful only for a unidirectional tree of messages in a single context. Engineering work needs to be done on identifiers that could be used in threads that cross media. Additionally, there is likely work to be done for messages that may not be strictly ordered in a thread.

The topics of "control panels" and automated introductions were deemed appropriate for further research.

#### 4.5. Message Transport

A central research topic that came out of the transport session was that of multiple transports. It was felt that much research could be done on the idea of transporting pieces of messages over separate transport media in order to get the message to its final destination. Especially in some high-latency, low-bandwidth environments, the ability to run parallel transports with different parts of messages could be extremely advantageous. The hard work in this area is re-associating all of the pieces in a timely manner, and identifying the single destination of the message when addressing will involve multiple media.

A common theme that arose in several of the discussions (including user control and message unification), but that figured prominently in the transport discussion, was a need for some sort of identifier. In the transport case, identifiers are necessary on two levels. Identifiers are needed to mark the endpoints in message transport. As described in the discussion, there are many cases where a message could reasonably be delivered to different entities that might all correspond to a single person. Some sort of identifier to indicate the target person of the message, as well as identifiers for the different endpoints, are all required in order to get any traction in this area. In addition, identifiers are also required for the messages being transported, as well as their component parts. Certainly, the idea of transporting different parts of a message over different mechanisms requires the identification of the containing message so that re-assembly can occur at the receiving end. However,

identifying the entire package is also necessary for those cases where duplicate copies of a message might be sent using two different mechanisms: The receiving end needs to find out that it has already received a copy of the message through one mechanism and identify that another copy of the message is simply a duplicate.

Workshop participants felt that, at the very least, a standard identifier syntax was a reasonable engineering work item that could be tackled. Though there exist some identifier mechanisms in current messaging protocols, none were designed to be used reliably across different transport environments or in multiple contexts. There is already a reasonable amount of engineering work done in the area of uniform resource identifiers (URI) that participants felt could be leveraged. Syntax would be required for identifiers of messages and their components as well as for identifiers of endpoint entities.

Work on the general problem of identifier use might have some tractable engineering aspects, especially in the area of message part identifiers, but workshop participants felt that more of the work was ripe for research. The ability to identify endpoints as belonging to a single recipient, and to be able to distribute identifiers of those endpoints with information about delivery preferences, is certainly an area where research could be fruitful. Additionally, it would be worthwhile to explore the collection of identified message components transported through different media, while delivering to the correct end-recipient with duplicate removal and re-assembly.

Package security was seen as an area for research. As described in Section 3.5, the possibility that different components of messages might travel over different media and need to be re-assembled at the recipient end breaks certain end-to-end security assumptions that are currently made. Participants felt that a worthwhile research goal would be to examine security mechanisms that could be used for such multi-component messages without sacrificing desirable security features.

Finally, a more architectural topic was that of restartability. Most current message transports, in the face of links with reliability problems, will cancel and restart the transport of a message from the beginning. Though some mechanisms do exist for restart mid-session, they are not widely implemented, and they certainly can rarely be used across protocol boundaries. Some architectural guidance on restart mechanisms would be a useful addition.

#### 4.6. Identity Hints and Key Distribution

It would be helpful to develop Internet-wide services to publish and retrieve keying material. One possible solution is to build such a service into Secure DNS, perhaps as an engineering item in an existing working group. However, care is needed since that would significantly increase the size and scope of DNS. A more research-oriented approach would be to investigate the feasibility of building Internet-wide key distribution services outside of DNS. In doing so, it is important to keep in mind that the problem of distribution is separate from the problem of enrollment, and that name subordination (control over what entities are allowed to create sub-domains) remains necessary.

Research may be needed to define the different audiences for message security. For example, users of consumer-oriented messaging services on the open Internet may not generally be willing or able to install new trusted roots in messaging client software, which may hamper the use of security technologies between businesses and customers. By contrast, within a single organization it may be possible to deploy new trusted roots more widely, since (theoretically) all of the organization's computing infrastructure is under the centralized control.

In defining security frameworks for messaging, it would be helpful to specify more clearly the similarities and differences among various messaging technologies with regard to trust models and messaging metaphors (e.g., stand-alone messages in email, discrete conversations in telephony, messaging sessions in instant messaging). The implications of these trust models and messaging metaphors for communications security have not been widely explored.

#### 5. Security Considerations

Security is discussed in several sections of this document, especially Sections 3.5, 3.6, 4.5, and 4.6.

#### 6. Acknowledgements

The IAB would like to thank QUALCOMM Incorporated for their sponsorship of the meeting rooms and refreshments.

The editors would like to thank all of the workshop participants. Eric Allman, Ted Hardie, and Cullen Jennings took helpful notes, which eased the task of writing this document.



## Appendix A. Participants

Eric Allman  
Nathaniel Borenstein  
Ben Campbell  
Dave Crocker  
Leslie Daigle  
Mark Day  
Mark Crispin  
Steve Dorner  
Lisa Dusseault  
Kevin Fall  
Ned Freed  
Randy Gellens  
Larry Greenfield  
Ted Hardie  
Joe Hildebrand  
Paul Hoffman  
Steve Hole  
Scott Hollenbeck  
Russ Housley  
Cullen Jennings  
Hisham Khartabil  
John Klensin  
John Levine  
Rohan Mahy  
Alexey Melnikov  
Jon Peterson  
Blake Ramsdell  
Pete Resnick  
Jonathan Rosenberg  
Peter Saint-Andre  
Greg Vaudreuil

## Appendix B. Pre-Workshop Papers

The topic papers circulated before the workshop were as follows:

Calendaring Integration (Nathaniel Borenstein)  
Channel Security (Russ Housley)  
Collaborative Authoring (Lisa Dusseault)  
Consent-Based Messaging (John Klensin)  
Content Security (Blake Ramsdell)  
Event Notifications (Joe Hildebrand)  
Extended Messaging Services (Dave Crocker)  
Group Messaging (Peter Saint-Andre)  
Identity and Reputation (John Levine)  
Instant Messaging and Presence Issues in Messaging (Ben Campbell)  
Large Email Environments (Eric Allman)  
Mail/News/Blog Convergence (Larry Greenfield)  
Messaging and Spam (Cullen Jennings)  
Messaging Metaphors (Ted Hardie)  
MUA/MDA, MUA/MSA, and MUA/Message-Store Interaction (Mark Crispin)  
Presence for Consent-Based Messaging (Jon Peterson)  
Rich Payloads (Steve Hole)  
Session-Oriented Messaging (Rohan Mahy)  
Spam Expectations for Mobile Devices (Greg Vaudreuil)  
Communication in Difficult-to-Reach Networks (Kevin Fall)  
Store-and-Forward Needs for IM (Hisham Khartabil)  
Syndication (Paul Hoffman)  
Transport Security (Alexey Melnikov)  
VoIP Peering and Messaging (Jonathan Rosenberg)  
Webmail, MMS, and Mobile Email (Randy Gellens)

## Authors' Addresses

Peter W. Resnick (Editor)  
Internet Architecture Board  
QUALCOMM Incorporated  
5775 Morehouse Drive  
San Diego, CA 92121-1714  
US

Phone: +1 858 651 4478  
EMail: [presnick@qualcomm.com](mailto:presnick@qualcomm.com)  
URI: <http://www.qualcomm.com/~presnick/>

Peter Saint-Andre (Editor)  
Jabber Software Foundation  
P.O. Box 1641  
Denver, CO 80201-1641  
US

Phone: +1 303 308 3282  
EMail: [stpeter@jabber.org](mailto:stpeter@jabber.org)  
URI: <http://www.jabber.org/people/stpeter.shtml>

## Full Copyright Statement

Copyright (C) The Internet Society (2006).

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

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

## Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

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

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

## Acknowledgement

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).