

Network Working Group
Request for Comments: 4473
Category: Informational

Y. Nomura
Fujitsu Labs
R. Walsh
J-P. Luoma
Nokia
J. Ott
Helsinki University of Technology
H. Schulzrinne
Columbia University
May 2006

Requirements for Internet Media Guides (IMGs)

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 memo specifies requirements for a framework and protocols for accessing and updating Internet Media Guide (IMG) information for media-on-demand and multicast applications. These requirements are designed to guide choice and development of IMG protocols for efficient and scalable delivery.

Table of Contents

1. Introduction	3
1.1. Background and Motivation	3
1.2. Scope of This Document	4
2. Terminology	5
2.1. New Terms	5
3. Problem Statement	6
4. Use Cases Requiring IMGs	7
4.1. Connectivity-based Use Cases	7
4.1.1. IP Datacast to a Wireless Receiver	7
4.1.2. Regular Fixed Dial-up Internet Connection	8
4.1.3. Broadband Always-on Fixed Internet Connection	9
4.2. Content-orientated Use Cases	9
4.2.1. TV and Radio Program Delivery	9
4.2.2. Media Coverage of a Live Event	10
4.2.3. Distance Learning	10
4.2.4. Multiplayer Gaming	10
4.2.5. File Distribution	11
4.2.6. Coming-release and Pre-released Content	11
5. Requirements	11
5.1. General Requirements	11
5.1.1. Independence of IMG Operations from IMG Metadata	11
5.1.2. Multiple IMG Senders	12
5.1.3. Modularity	12
5.2. Delivery Properties	12
5.2.1. Scalability	13
5.2.2. Support for Intermittent Connectivity	13
5.2.3. Congestion Control	13
5.2.4. Sender- and Receiver-Driven Delivery	13
5.3. Customized IMGs	14
5.4. Reliability	15
5.4.1. Managing Consistency	15
5.4.2. Reliable Message Exchange	16
5.5. IMG Descriptions	16
6. Security Considerations	17
6.1. IMG Authentication and Integrity	18
6.2. Privacy	19
6.3. Access Control for IMGs	19
6.4. Denial-of-Service (DOS) Attacks	20
6.5. Replay Attacks	20
7. Normative References	21
8. Informative References	21
9. Acknowledgements	22

1. Introduction

1.1. Background and Motivation

For some ten years, multicast-based (multimedia) conferences (including IETF working group sessions) as well as broadcasts of lectures/seminars, concerts, and other events have been used in the Internet, more precisely, on the MBONE. Schedules and descriptions for such multimedia sessions as well as the transport addresses, codecs, and their parameters have been described using the Session Description Protocol (SDP) [2] as a rudimentary (but as of then largely sufficient) means. Descriptions have been disseminated using the Session Announcement Protocol (SAP) [3] and Session Directory Tools such as SD [4] or SDR [5]; descriptions have also been put up on web pages, sent by electronic mail, etc.

Recently, interest has grown to expand -- or better: to generalize -- the applicability of these kinds of session descriptions. Descriptions are becoming more elaborate in terms of included metadata, more generic regarding the types of media sessions, and possibly also support other transports than just IP (e.g., legacy TV channel addresses). This peers well with the DVB (Digital Video Broadcasting) [6] Organization's increased activities towards IP-based communications over satellite, cable, and terrestrial radio networks, also considering IP as the basis for TV broadcasts and further services. The program/content descriptions are referred to as Internet Media Guides (IMGs) and can be viewed as a generalization of Electronic Program Guides (EPGs) and multimedia session descriptions.

An Internet Media Guide (IMG) has a structured collection of multimedia session descriptions expressed using SDP, SDPng [7], or some similar session description format. This is used to describe a set of multimedia services (e.g., television program schedules, content delivery schedules) but may also refer to other networked resources including web pages. IMGs provide the envelope for metadata formats and session descriptions defined elsewhere with the aim of facilitating structuring, versioning, referencing, distributing, and maintaining (caching, updating) such information.

The IMG metadata may be delivered to a potentially large audience, who uses it to join a subset of the sessions described, and who may need to be notified of changes to this information. Hence, a framework for distributing IMG metadata in various different ways is needed to accommodate the needs of different audiences: For traditional broadcast-style scenarios, multicast-based (push) distribution of IMG metadata needs to be supported. Where no multicast is available, unicast-based push is required, too.

Furthermore, IMG metadata may need to be retrieved interactively, similar to web pages (e.g., after rebooting a system or when a user is browsing after network connectivity has been re-established). Finally, IMG metadata may be updated as time elapses because content described in the guide may be changed: for example, the airtime of an event such as a concert or sports event may change, possibly affecting the airtime of subsequent media. This may be done by polling the IMG sender as well as by asynchronous change notifications.

Furthermore, any Internet host can be a sender of content and thus an IMG sender. Some of the content sources and sinks may only be connected to the Internet sporadically. Also, a single human user may use many different devices to access metadata. Thus, we envision that IMG metadata can be sent and received by, among others, cellular phones, Personal Digital Assistants (PDAs), personal computers, streaming video servers, set-top boxes, video cameras, and Digital Video Recorders (DVRs), and that the data be carried across arbitrary types of link layers, including bandwidth-constrained mobile networks. However, generally we expect IMG senders to be well-connected hosts.

Finally, with many potential senders and receivers, different types of networks, and presumably numerous service providers, IMG metadata may need to be combined, split, filtered, augmented, modified, etc., on their way from the sender(s) to the receiver(s) to provide the ultimate user with a suitable selection of multimedia services according to her preferences, subscriptions, location, and context (e.g., devices, access networks).

1.2. Scope of This Document

This document defines requirements that Internet Media Guide mechanisms must satisfy in order to deliver IMG metadata to a potentially large audience. Since IMGs can describe many kinds of multimedia content, IMG methods are generally applicable to several scenarios.

In considering wide applicability, this document provides the problem statement and discusses existing mechanisms in this area. Then several use case scenarios for IMGs are explained including descriptions of how IMG metadata and IMG delivery mechanisms contribute to these scenarios. Following this, this document provides general requirements that are independent of any transport layer mechanism and application, such as delivery properties, reliability, and IMG descriptions.

This document reflects investigating work on delivery mechanisms for IMGs and generalizing work on session announcement and initiation protocols, especially in the field of the MMUSIC working group (SAP, SIP [8], SDP).

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [1].

2.1. New Terms

Internet Media Guide (IMG): IMG is a generic term used to describe the formation, delivery, and use of IMG metadata. The definition of the IMG is intentionally left imprecise.

IMG Element: The smallest atomic element of metadata that can be transmitted separately by IMG operations and referenced individually from other IMG elements.

IMG Metadata: A set of metadata consisting of one or more IMG elements. IMG metadata describes the features of multimedia content used to enable selection of and access to media sessions containing content. For example, metadata may consist of the URI, title, airtime, bandwidth needed, file size, text summary, genre, and access restrictions.

IMG Delivery: The process of exchanging IMG metadata in terms of both large-scale and atomic data transfers.

IMG Sender: An IMG sender is a logical entity that sends IMG metadata to one or more IMG receivers.

IMG Receiver: An IMG receiver is a logical entity that receives IMG metadata from an IMG sender.

IMG Transceiver: An IMG transceiver combines an IMG receiver and sender. It may modify received IMG metadata or merge IMG metadata received from several different IMG senders.

IMG Operation: An atomic operation of an IMG transport protocol, used between IMG sender(s) and IMG receiver(s) for the delivery of IMG metadata and for the control of IMG sender(s)/receiver(s).

IMG Transport Protocol: A protocol that transports IMG metadata from an IMG sender to IMG receiver(s).

IMG Transport Session: An association between an IMG sender and one or more IMG receivers within the scope of an IMG transport protocol. An IMG transport session involves a time-bound series of IMG transport protocol interactions that provide delivery of IMG metadata from the IMG sender to the IMG receiver(s).

3. Problem Statement

As we enumerate the requirements for IMGs, it will become clear that they are not fully addressed by the existing protocols. The "Framework for the Usage of Internet Media Guides" [9] discusses about these issues in more detail.

The MMUSIC working group has long been investigating content, media and service information delivery mechanisms, and protocols, and has itself produced the Session Announcement Protocol (SAP), the Session Description Protocol (SDP), and the Session Initiation Protocol (SIP). SDP is capable of describing multimedia sessions (i.e., content in a wider sense) by means of limited descriptive information intended for human perception plus transport, scheduling information, and codecs and addresses for setting up media sessions. SIP and SAP are protocols to distribute these session descriptions.

However, we perceive a lack of a standard solution for scalable IMG delivery mechanism in the number of receivers with consistency of IMG metadata between an IMG sender and IMG receiver for both bi-directional and unidirectional transport. With increased service dynamics and complexity, there is an increased requirement for updates to these content descriptions.

HTTP [10] is a well-known information retrieval protocol using bi-directional transport and is widely used to deliver web-based content descriptions to many hosts. However, it has well-recognized limitations of scalability in the number of HTTP clients since it relies on the polling mechanism to keep information consistent between the server and client.

SAP [3] is an announcement protocol that distributes session descriptions via multicast. It does not support prioritization or fine-grained metadata selection and update notifications, as it places restrictions on metadata payload size and always sends the whole metadata. It requires a wide-area multicast infrastructure for it to be deployable beyond a local area network.

SIP [8] and SIP-specific event notifications [11] can be used to notify subscribers of the update of IMG metadata for bi-directional transport. However, it is necessary to define an event package for IMGs.

We also perceive a lack of standard solution for flexible content descriptions to support a multitude of application-specific metadata and associated data models with a different amount of detail and different target audiences.

SDP [2] has a text-encoded syntax to specify multimedia sessions for announcements and invitations. It is primarily intended to describe client capability requirements and enable client application selection. Although SDP is extensible, it has limitations such as structured extensibility and capability to reference properties of a multimedia session from the description of another session.

These can mostly be overcome by the XML-based SDPng [7] -- which is intended for both two-way negotiation and unidirectional delivery -- or similar content description mechanisms. Since SDPng addresses multiparty multimedia conferences, it would be necessary to extend the XML schema in order to describe general multimedia content.

4. Use Cases Requiring IMGs

4.1. Connectivity-based Use Cases

4.1.1. IP Datacast to a Wireless Receiver

IP Datacast is the delivery of IP-based services over broadcast radio. Internet content delivery is therefore unidirectional in this case. However, there can be significant benefits from being able to provide rich media one-to-many services to such receivers.

There are two main classes of receiver in this use case: fixed mains-powered and mobile battery-powered. Both of these are affected by radio phenomena and so robust, or error-resilient, delivery is important. Carouselled metadata transfer (cyclically repeated with a fixed bandwidth) provides a base level of robustness for an IP datacast-based announcement system, although the design of carouselled transfer should enable battery-powered receivers to go through periods of sleep to extend their operational time between charges. Insertion of Forward Error Correction (FEC) data into metadata announcements improves error resilience, and reordering (interleaving) data blocks further increases tolerance to burst errors.

To enable receivers to more accurately specify the metadata they are interested in, the unidirectional delivery may be distributed between several logical channels. This is so that a receiver needs only access the channels of interest and thus can reduce the amount of time, storage, and CPU resources needed for processing the IP data. Also, hierarchical channels enable receivers to subscribe to a (possibly well-known) root multicast channel/group and progressively access only those additional channels based on metadata in parent channels.

In some cases, the receiver may have multiple access interfaces adding bi-directional communications capability. This enables a multitude of options, but most important, it enables NACK-based reliability and the individual retrieval of missed or not-multicast sets of metadata.

Thus, essential IMG features in this case include the following: robust unidirectional delivery (with optional levels of reliability including "plug-in FEC" supported by a transport layer protocol), which implies easily identifiable segmentation of delivery data to make FEC, carousel, interleaving, and other schemes possible; effective identification of metadata sets (probably uniquely) to enable more efficient use of multicast and unicast retrieval over multiple access systems regardless of the parts of metadata and application-specific extensions in use; and prioritization of metadata, which can (for instance) be achieved by spreading it between channels and allocating/distributing bandwidth accordingly.

Furthermore, some cases require IMG metadata authentication and some group security/encryption and supporting security message exchanges (out of band from the IMG multicast sessions).

4.1.2. Regular Fixed Dial-up Internet Connection

Dial-up connections tend to be reasonably slow (<56 kbps in any case), and thus large data transfers are less feasible, especially during an active application session (such as a file transfer described by IMG metadata). They can also be intermittent, especially if a user is paying for the connected time, or connected through a less reliable exchange. Thus, this favors locally stored IMG metadata over web-based browsing, especially where parts of the metadata change infrequently. There may be no service provider preference over unicast and multicast transport for small and medium numbers of users as the last-mile dial-up connection limits per-user congestion, and a user may prefer the more reliable option (unicast unless reliable multicast is provided).

4.1.3. Broadband Always-on Fixed Internet Connection

Typically, bandwidth is less of an issue to a broadband user and unicast transport, such as using query-response methods, may be typical for a PC user. If a system were only used in this context, with content providers, ISPs, and users having no other requirements, then web-based browsing may be equally suitable. However, broadband users sharing a local area network, especially wireless, may benefit more from local storage features than on-line browsing, especially if they have intermittent Internet access.

Some services on broadband, such as live media broadcasting, benefit from multicast transport for streaming media because of scalability. In the cases where multicast transport is already available, it is convenient for a sender and receiver to retrieve IMG metadata over multicast transport. Thus, broadband users may be forced to retrieve IMG metadata over multicast if backbone operators require this to keep system-wide bandwidth usage feasible.

4.2. Content-orientated Use Cases

IMGs will be able to support a very wide range of use cases for enabling content/media delivery. The following few sections just touch the surface of what is possible and are intended to provide an understanding of the scope and type of IMG usage. Many more examples may be relevant, for instance, those detailed in [12]. There are several unique features of IMGs that set them apart from related application areas such as Service Location Protocol (SLP) based service location discovery, Lightweight Directory Access Protocol (LDAP) based indexing services, and search engines such as Google. Features unique to IMGs include the following:

- o IMG metadata is generally time-related
- o There are timeliness requirements in the delivery of IMG metadata
- o IMG metadata may be updated as time elapses or when an event arises

4.2.1. TV and Radio Program Delivery

A sender of audio/video streaming content can use the IMG metadata to describe the scheduling and other properties of audio/video sessions and events within those sessions, such as individual TV and radio programs and segments within those programs. IMG metadata describing audio/video streaming content could be represented in a format

similar to that of a TV guide in newspapers, or an Electronic Program Guide available on digital TV receivers.

TV and radio programs can be selected for reception either manually by the end-user or automatically based on the content descriptions and the preferences of the user. The received TV and radio content can be either presented in real time or recorded for later consumption. There may be changes in the scheduling of a TV or a radio program, possibly affecting the transmission times of subsequent programs. IMG metadata can be used to notify receivers of such changes, enabling users to be prompted or recording times to be adjusted.

4.2.2. Media Coverage of a Live Event

The media coverage of a live event such as a rock concert or a sports event is a special case of regular TV/radio programming. There may be unexpected changes in the scheduling of a live event, or the event may be unscheduled to start with (such as breaking news). In addition to audio/video streams, textual information relevant to the event (e.g., statistics of the players during a football match) may be sent to user terminals. Different transport modes or even different access technologies can be used for the different media: for example, a unidirectional datacast transport could be used for the audio/video streams and an interactive cellular connection for the textual data. IMG metadata should enable terminals to discover the availability of different media used to cover a live event.

4.2.3. Distance Learning

IMG metadata could describe compound sessions or services enabling several alternative interaction modes between the participants. For example, the combination of one-to-many media streaming, unicast messaging, and downloading of presentation material could be useful for distance learning.

4.2.4. Multiplayer Gaming

Multiplayer games are an example of real-time multiparty communication sessions that could be advertised using IMGs. A gaming session could be advertised either by a dedicated server or by the terminals of individual users. A user could use IMGs to learn of active multiplayer gaming sessions, or advertise the user's interest in establishing such a session.

4.2.5. File Distribution

IMGs support the communication of file delivery session properties, enabling the scheduled delivery or synchronization of files between a number of hosts. The received IMG metadata could be subsequently used by any application (also outside the scope of IMGs), for example, to download a file with a software update. IMG metadata can provide a description to each file in a file delivery session, assisting users or receiving software in selecting individual files for reception.

For example, when a content provider wants to distribute a large amount of data in file format to thousands of clients, the content provider can use IMG metadata to schedule the delivery effectively.

Since IMG metadata can describe time-related data for each receiver, the content provider can schedule delivery time for each receiver. This can save network bandwidth and delivery capacity of senders. In addition, IMG metadata can be used to consistency check, and thus synchronize, a set of files between a sender host and receiver host, when those files change as time elapses.

4.2.6. Coming-release and Pre-released Content

IMG metadata can be used to describe items of content before the details of their final release are known. A user may be interested in coming content (a new movie or software title where some aspects of the content description are known in advance) and so subscribe to an information service that notifies the user of changes to metadata describing that content. Thus, as the coming release (or pre-releases, e.g., as movie trailers or software demos) become available, the IMG metadata changes and the user is notified of this change. For example, the user could see an announcement of a movie that will be released sometime in the next few months, and configure the user's terminal to receive and record any trailers or promotional material as they become available.

5. Requirements

5.1. General Requirements

5.1.1. Independence of IMG Operations from IMG Metadata

REQ GEN-1: Carrying different kinds of IMG metadata format and different IMG metadata formats in the IMG message body MUST be allowed.

REQ GEN-2: Delivery mechanisms SHOULD support many different applications' specific metadata formats to keep the system interoperable with existing applications.

This provides flexibility in selecting/designing an IMG transport protocol suited to various scenarios.

5.1.2. Multiple IMG Senders

REQ GEN-3: IMG receivers MUST be allowed to communicate with any number of IMG senders simultaneously.

This might lead to receiving redundant IMG metadata describing the same items; however, it enables the IMG receiver access to more IMG metadata than may be available from a single IMG sender. This also provides flexibility for the IMG transport protocols and does not preclude a mechanism to solve inconsistency among IMG metadata due to multiple IMG senders. This document assumes that a typical IMG environment will involve many more IMG receivers than IMG senders and that IMG senders are continually connected for the duration of interest (rather than intermittently connected).

5.1.3. Modularity

REQ GEN-4: The IMG delivery mechanisms MUST allow the combination of several IMG operations.

This is for the purpose of extending functionality (e.g., several or one protocol(s) to provide all the needed operations). Applications can select an appropriate operation set to fulfill their purpose.

5.2. Delivery Properties

This section describes general performance requirements based on the assumption that the range of IMG usage shall be important. However, note that requirements for delivery properties may vary based on the usage scenario, and thus some limited-use implementations place less importance on some requirements.

For example, it is clear that a multicast transport may provide more scalable delivery than a unicast transport; however, scalability requirements do not preclude the unicast transport mechanisms. In this sense, scalability is always important for the protocols irrespective of transport mechanisms.

5.2.1. Scalability

REQ DEL-1: The IMG system MUST be scalable to large numbers of messages, so as to allow design and use of delivery mechanisms that will not fail in delivering up-to-date information under huge numbers of transactions and massive quantities of IMG metadata.

REQ DEL-2: IMGs SHOULD provide a method to prevent an IMG sender from sending unnecessary IMG metadata that have been stored or deleted in IMG receivers.

REQ DEL-3: The protocol MUST be scalable to very large audience sizes requiring IMG delivery.

5.2.2. Support for Intermittent Connectivity

REQ DEL-4: The system MUST enable IMG receivers with intermittent access to network resources (connectivity) to receive and adequately maintain sufficient IMG metadata.

This allows intermittent access to save power where there is no need to keep communications links powered up while they are sitting idle. For instance, in this situation, periodic bursts of notifies or a fast cycling update carousel allow hosts to wake up for short periods of time and still be kept up-to-date. This can be beneficial for IMG receivers with sporadic connections to the fixed Internet, but is critical in the battery-powered wireless Internet.

The implication of intermittent connectivity is that immediate distribution of changes becomes infeasible and so managing data consistency should be focused on the timely delivery of data.

5.2.3. Congestion Control

REQ DEL-5: Internet-friendly congestion control MUST be provided for use on the public Internet.

REQ DEL-6: An IMG entity SHOULD invalidate the IMG metadata item when an IMG metadata item has lifetime information and its lifetime is over. This will lessen the need for notifications of updates from the IMG sender to the IMG receiver to invalidate the item and may help in reducing load.

5.2.4. Sender- and Receiver-Driven Delivery

REQ DEL-7: The system MUST be flexible in choosing sender-driven, receiver-driven, or both delivery schemes.

Sender-driven delivery achieves high scalability without interaction between the IMG sender and receiver.

In contrast, receiver-driven delivery provides on-demand delivery for IMG receivers. Since an IMG sender's complete IMG metadata may be a very large amount of data, the IMG receiver needs to be able to access the guide when convenient (e.g., when sufficient network bandwidth is available to the IMG receiver).

5.3. Customized IMGs

REQ CUS-1: The system MUST allow delivery of customized IMG metadata.

The IMG receiver may require a subset of all the IMG metadata available according to their preferences (type of content, media description, appropriate age group, etc.) and configuration.

The IMG receiver might send its preferences in the IMG operations that can specify user-specific IMG metadata to be delivered. These preferences could consist of filtering rules. When receiving these messages, the IMG sender might respond with appropriate messages carrying a subset of IMG metadata that matches the IMG receiver's preferences.

This mechanism can reduce the amount of IMG metadata delivered from the IMG sender to IMG receiver, and consequently it can save the resource consumption on the IMG entities and networks. It is typically useful in unicast cases and also beneficial in multicast cases where an IMG sender distributes the same IMG metadata to interested IMG receivers at the same time.

For multicast and unicast cases where the IMG sender does not provide customized IMG metadata, the IMG receiver could receive all IMG metadata transmitted on the channels that the IMG receiver has joined. However, it may select and filter the IMG metadata to get customized IMG metadata by its preferences, and thus drop unwanted metadata immediately upon reception.

Customizing metadata might be achieved by changing the IMG descriptions sent and IMG receivers and/or changing the delivery properties (channels used).

Note that customization and scalability are only somewhat exclusive. Systems providing an IMG receiver to an IMG sender request-based customization will be generally less scalable to massive IMG receiver populations than those without this return signaling technique. Thus, customization, as with any feature that affects scalability, should be carefully designed for the intended application, and it may

not be possible that a one-size-fits-all solution for customization would meet the scalability requirements for all applications and deployment cases.

5.4. Reliability

5.4.1. Managing Consistency

IMG metadata tends to change as time elapses; as new content is added, the old IMG metadata stored in the IMG receiver becomes unavailable, and the parameters of the existing IMG metadata are changed.

REQ REL-1: The system MUST manage IMG metadata consistency.

Either the IMG sender can simply make updates available (unsynchronized), or the IMG sender and receiver can interact to keep their copies of the IMG metadata synchronized.

In the unsynchronized model, the IMG sender does not know whether a particular IMG receiver has an up-to-date copy of the IMG metadata.

In the synchronized model, updating a cached copy of the IMG metadata is necessary to control consistency when the IMG sender or receiver could not communicate for a while. In this case, the IMG sender or receiver may need to confirm its consistency by IMG operations.

REQ REL-2: Since IMG metadata can change at any time, IMG receivers SHOULD be notified of such changes.

Fulfilling this requirement needs to be compatible with the scalability requirements for the number of IMG receivers and the consistency of metadata.

Depending on the size of the IMG metadata, the interested party may want to defer retrieving the actual information. The change notification should be addressed to a logical user (or user group), rather than a host, since users may change devices.

Note that depending on the deployment environment and application specifics, the level of acceptable inconsistency varies. Thus, this document does not define inconsistency as specific time and state differences between IMG metadata stored in an IMG sender and IMG metadata stored in an IMG receiver.

In general, the consistency of metadata for content and media is more important immediately prior to and during the media's session(s). Hosts that forward (or otherwise resend) metadata may not tolerate

inconsistencies because delivering out-of-date data is both misleading and bandwidth inefficient.

5.4.2. Reliable Message Exchange

REQ REL-4: An IMG transport protocol MUST support reliable message exchange.

The extent to which this could result in 100% error-free delivery to 100% of IMG receivers is a statistical characteristic of the protocols used. Usage of reliable IMG delivery mechanisms is expected to depend on the extent to which underlying networks provide reliability and, conversely, introduce errors. Note that some deployments of IMG transport protocols may not aim to provide perfect reception to all IMG receivers in all possible cases.

5.5. IMG Descriptions

REQ DES-1: IMG metadata MUST be interoperable over any IMG transport protocol, such that an application receiving the same metadata over any one (or more) of several network connections and/or IMG transport protocols will interpret the metadata in exactly the same way. (This also relates to the 'Independence of IMG Operations from IMG Metadata' requirements.)

REQ DES-2: IMG delivery MUST enable the carriage of any format of application-specific metadata.

Thus, the system will support the description of many kinds of multimedia content, without the need for a single homogeneous metadata syntax for all uses (which would be infeasible anyway). This is essential for environments using IMG systems to support many kinds of multimedia content and to achieve wide applicability.

REQ DES-3: Whereas specific applications relying on IMGs will need to select one or more specific application-specific metadata formats (standard, syntax, etc.), the IMG system MUST be independent of this (it may be aware, but it will operate in the same way for all).

Thus, a metadata transfer envelope format that is uniform across all different application-specific IMG metadata formats is needed. The envelope would reference (point to) or carry (as payload) some application-specific metadata, and the envelope would support the maintenance of the application-specific metadata, which may also serve the metadata relationships determined by the data model(s) used. The envelope would not need to be aware of the data model(s) in use.

REQ DES-4: IMG metadata MUST be structured to enable fragmentation for efficient delivery.

This is intended to ensure that an IMG sender with more than a trivial knowledge of metadata is able to deliver only part of its (and the global) complete IMG metadata knowledge. (For instance, a trivial quantity of knowledge could be a single SDP description.) In general, the resolution of this fragmentation will be very much dependent on the optimal delivery of a deployment, although some metadata syntaxes will inherently affect the sensible lower limit for a single element/fragment.

REQ DES-5: A metadata transfer envelope MUST be defined to include essential parameters.

Examples of essential parameters are those that allow the metadata in question to be uniquely identified and updated by new versions of the same metadata.

REQ DES-6: It SHALL be possible to deduce the metadata format via the metadata transfer envelope.

REQ DES-7: IMG senders SHALL use a metadata transfer envelope for each IMG metadata transfer.

Thus, it will even be possible to describe relationships between syntactically dissimilar application-specific formats within the same body of IMG metadata knowledge. (For instance, a data model could be instantiated using both SDP and SDPng.)

REQ DES-8: IMG metadata SHOULD support the description of differences between an updated version and an old version of IMG metadata when the IMG delivery mechanism carries updated IMG metadata and those differences are considerably little (e.g., by providing a 'delta' of the two versions; this also relates the delivery property requirements for congestion control in Section 5.2.3).

6. Security Considerations

Internet Media Guides are used to convey information about multimedia resources from one or more IMG senders across one or more intermediaries to one or more IMG receivers. IMG metadata may be pushed to the IMG receivers or interactively retrieved by them. IMGs provide metadata as well as scheduling and rendezvous information about multimedia resources, and so on, and requests for IMG metadata may contain information about the requesting users.

The information contained in IMG metadata as well as the operations related to IMGs should be secured to avoid forged information, misdirected users, and spoofed IMG senders, for example, and to protect user privacy.

The remainder of this section addresses the security requirements for IMGs.

6.1. IMG Authentication and Integrity

IMG metadata and its parts need to be protected against unauthorized alteration/addition/deletion on the way. Their originator needs to be authenticated.

REQ AUT-1: It MUST be possible to authenticate the originator of a set of IMG metadata.

REQ AUT-2: It MUST be possible to authenticate the originator of a subpart of IMG metadata (e.g., a delta or a subset of the information).

REQ AUT-3: It MUST be possible to validate the integrity of IMG metadata.

REQ AUT-4: It MUST be possible to validate the integrity of a subpart of IMG metadata (e.g., a delta or a subset of the information).

REQ AUT-5: It MUST be possible to separate or combine individually authenticated pieces of IMG metadata (e.g., in an IMG transceiver) without invalidating the authentication.

REQ AUT-6: It MUST be possible to validate the integrity of an individually authenticated piece of IMG metadata even after this piece has been separated from other pieces of IMG metadata and combined with other pieces to form new IMG metadata.

REQ AUT-7: It MUST be possible to authenticate the originator of an IMG operation.

REQ AUT-8: It MUST be possible to validate the integrity of any contents of an IMG operation (e.g., the subscription or inquiry information).

6.2. Privacy

Customized IMG metadata and IMG metadata delivered by notification to individual users may reveal information about the habits and preferences of a user and may thus deserve confidentiality protection, even though the information itself is public.

REQ PRI-1: It MUST be possible to keep user requests to subscribe to or retrieve certain (parts of) IMG metadata confidential.

REQ PRI-2: It MUST be possible to keep IMG metadata, pieces of IMG metadata, or pointers to IMG metadata delivered to individual users or groups of users confidential.

REQ PRI-3: It SHOULD be possible to ensure this confidentiality end-to-end, that is, to prevent intermediaries (such as IMG transceivers) from accessing the contained information.

6.3. Access Control for IMGs

Some IMG metadata may be freely available, while access to other IMG metadata may be restricted to closed user groups (e.g., paying subscribers). Also, different parts of IMG metadata may be protected at different levels: for example, metadata describing a media session may be freely accessible, while rendezvous information to actually access the media session may require authorization.

REQ ACC-1: It MUST be possible to authorize user access to IMG metadata.

REQ ACC-2: It MUST be possible to authorize access of users to pieces of IMG metadata (delta information, subparts, pointers).

REQ ACC-3: It MUST be possible to require different authorization for different parts of the same IMG metadata.

REQ ACC-4: It MUST be possible to access selected IMG metadata anonymously.

REQ ACC-5: It MUST be possible for an IMG receiver to choose not to receive (parts of) IMG metadata in order to avoid being identified by the IMG sender.

REQ ACC-6: It SHOULD be possible for an IMG transceiver to select suitable authorization methods that are compatible between both IMG senders and IMG receivers it interacts with.

REQ ACC-7: It MAY be possible for IMG senders to require certain authorization that cannot be modified by intermediaries.

6.4. Denial-of-Service (DOS) Attacks

Retrieving or distributing IMG metadata may require state in the IMG senders, transceivers, and/or receivers for the respective IMG transport sessions. Attackers may create large numbers of sessions with any of the above IMG entities to disrupt regular operation.

REQ DOS-1: IMG operations SHOULD be authenticated.

REQ DOS-2: It SHOULD be possible to avoid DoS attacks that build up session state in IMG entities to exhaust their resources.

REQ DOS-3: It SHOULD be possible to avoid DoS attacks that exhaust resources of IMG entities by flooding them with IMG metadata.

As an example, two potential solutions that may be considered are running an IMG entity in stateless mode or identification and discarding of malicious packets by an IMG entity.

6.5. Replay Attacks

IMG metadata disseminated by an IMG sender or an IMG transceiver may be updated, be deleted, or lose validity over time for some other reasons. Replaying outdated IMG metadata needs to be prevented.

Furthermore, replay attacks may also apply to IMG operations (rather than just their payload). Replaying operations also needs to be prevented.

REQ REP-1: IMG metadata MUST be protected against partial or full replacement of newer ("current") versions by older ones.

In a system with multiple senders, it may not be feasible to prevent some senders from delivering older versions of metadata than others - as a result of imperfect sender-sender data consistency. Thus, replay attacks and delivery of inconsistent data require that an IMG receiver verifies that the IMG metadata is valid and reliable by using some security mechanism(s) (e.g., authorization, authentication, or integrity).

REQ REP-2: Mechanisms MUST be provided to mitigate replay attacks on the IMG operations.

The level of threat from replay attacks varies very much depending on system scale and how well defined or open it is. Thus, mitigating replay attacks may lead to different solutions for different systems, independent of the basic delivery method and metadata definitions. A system with multiple senders presents a more challenging scenario for handling replay attacks. As an example, bundling metadata with a security mechanism is one potential solution.

7. Normative References

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

8. Informative References

- [2] Handley, M. and V. Jacobson, "SDP: Session Description Protocol", RFC 2327, April 1998.
- [3] Handley, M., Perkins, C., and E. Whelan, "Session Announcement Protocol", RFC 2974, October 2000.
- [4] Session Directory, <ftp://ftp.ee.lbl.gov/conferencing/sd/>
- [5] Session Directory Tool, <http://www-mice.cs.ucl.ac.uk/multimedia/software/sdr/>
- [6] Digital Video Broadcasting Project, <http://www.dvb.org/>
- [7] Kutscher, D., Ott, J., and C. Bormann, "Session description and capability negotiation", Work in Progress, February 2005.
- [8] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.
- [9] Nomura, Y., Walsh, R., Luoma, J-P., Asaeda, H., and H. Schulzrinne, "Framework for the Usage of Internet Media Guides (IMG)", RFC 4435, April 2006.
- [10] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999.
- [11] Roach, A.B., "Session Initiation Protocol (SIP)-Specific Event Notification", RFC 3265, June 2002.
- [12] Quinn, B. and K. Almeroth, "IP Multicast Applications: Challenges and Solutions", RFC 3170, September 2001.

9. Acknowledgements

The authors would like to thank Hitoshi Asaeda, Gonzalo Camarillo, Jean-Pierre Evain, Dirk Kutscher, Petri Koskelainen, Colin Perkins, Toni Paila, and Magnus Westerlund for their excellent comments and ideas on this work.

Authors' Addresses

Yuji Nomura
Fujitsu Laboratories Ltd.
4-1-1 Kamikodanaka, Nakahara-ku, Kawasaki 211-8588
Japan

EEmail: nom@flab.fujitsu.co.jp

Rod Walsh
Nokia Research Center
P.O. Box 100, FIN-33721 Tampere
Finland

EEmail: rod.walsh@nokia.com

Juha-Pekka Luoma
Nokia Research Center
P.O. Box 100, FIN-33721 Tampere
Finland

EEmail: juha-pekka.luoma@nokia.com

Joerg Ott
Helsinki University of Technology
Networking Laboratory
PO Box 3000
FIN-02015 TKK
Finland

EEmail: jo@netlab.tkk.fi

Henning Schulzrinne
Dept. of Computer Science
Columbia University
1214 Amsterdam Avenue
New York, NY 10027
USA

EEmail: schulzrinne@cs.columbia.edu

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.

Acknowledgement

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