

Network Working Group
Request for Comments: 4435
Category: Informational

Y. Nomura
Fujitsu Labs.
R. Walsh
J-P. Luoma
Nokia
H. Asaeda
Keio University
H. Schulzrinne
Columbia University
April 2006

A Framework for the Usage of 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 document defines a framework for the delivery of Internet Media Guides (IMGs). An IMG is a structured collection of multimedia session descriptions expressed using the Session Description Protocol (SDP), SDPng, or some similar session description format. This document describes a generalized model for IMG delivery mechanisms, the use of existing protocols, and the need for additional work to create an IMG delivery infrastructure.

Table of Contents

1.	Introduction	3
2.	Terminology	3
2.1.	New Terms	4
3.	IMG Common Framework Model	5
3.1.	IMG Data Types	5
3.1.1.	Complete IMG Description	5
3.1.2.	Delta IMG Description	6
3.1.3.	IMG Pointer	6
3.2.	IMG Entities	6
3.3.	Operation Set for IMG Delivery	7
3.3.1.	IMG ANNOUNCE	7
3.3.2.	IMG QUERY	8
3.3.3.	IMG RESOLVE	8
3.3.4.	IMG SUBSCRIBE	8
3.3.5.	IMG NOTIFY	9
3.3.6.	Binding between IMG Operations and Data Types	9
3.4.	Overview of Protocol Operations	10
4.	Deployment Scenarios for IMG Entities	10
4.1.	Intermediary Cases	10
4.2.	One-to-Many Unidirectional Multicast	12
4.3.	One-to-One Bidirectional Unicast	12
4.4.	Combined Operations with Common Metadata	13
5.	Applicability of Existing Protocols to the Proposed Framework Model	14
5.1.	Existing Standard Fitting the IMG Framework Model	14
5.2.	IMG Mechanism Needs Which Are Not Met by Existing Standards	15
5.2.1.	A Multicast Transport Protocol	16
5.2.2.	Usage of Unicast Transport Protocols	16
5.2.3.	IMG Envelope	17
5.2.4.	Metadata Data Model	18
6.	Security Considerations	18
7.	Normative References	19
8.	Informative References	19
9.	Acknowledgements	20

1. Introduction

Internet Media Guides (IMGs) provide and deliver structured collections of multimedia descriptions expressed using SDP [2], SDPng [3], or other description formats. They are used to describe sets of multimedia services (e.g., television program schedules, content delivery schedules) and refer to other networked resources including web pages. IMGs provide an envelope for metadata formats and session descriptions defined elsewhere with the aim of facilitating structuring, versioning, referencing, distributing, and maintaining (caching, updating) such information.

IMG metadata may be delivered to a potentially large audience, which uses it to join a subset of the sessions described and which may need to be notified of changes to the IMG metadata. 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.

This document defines a common framework model for IMG delivery mechanisms and their deployment in network entities. There are three fundamental components in the IMG framework model: data types, operation sets, and entities. These components specify a set of framework guidelines for efficient delivery and description of IMG metadata. The data types give generalized means to deliver and manage the consistency of application-specific IMG metadata. IMG operations cover broadcast, multicast distribution, event notification upon change, unicast-based push, and interactive retrievals similar to web pages.

Since we envision that any Internet host can be a sender and receiver of IMG metadata, a host involved in IMG operations performs one or more of the roles defined as the entities in the IMG framework model. The requirements for IMG delivery mechanisms and descriptions can be found in the IMG requirements document [4].

This document outlines the use of existing protocols to create an IMG delivery infrastructure. It aims to organize existing protocols into a common model and show their capabilities and limitations from the viewpoint of IMG delivery functions.

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 to describe the formation, delivery, and use of IMG metadata. The definition of the IMG is intentionally left imprecise [4].

IMG Element: The smallest atomic element of metadata that can be transmitted separately by IMG operations and referenced individually from other IMG elements [4].

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 a media object's URI, title, airtime, bandwidth needed, file size, text summary, genre, and access restrictions [4].

IMG Description: A collection of IMG metadata with a data type indicating a self-contained set or a subset of IMG metadata, or a reference to IMG metadata.

IMG Delivery: The process of exchanging IMG metadata both in terms of large-scale and atomic data transfers [4].

IMG Sender: An IMG sender is a logical entity that sends IMG metadata to one or more IMG receivers [4].

IMG Receiver: An IMG receiver is a logical entity that receives IMG metadata from an IMG sender [4].

IMG Transceiver: An IMG transceiver combines an IMG receiver and sender. It may modify received IMG metadata or merge IMG metadata received from a several different IMG senders [4].

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) [4].

IMG Transport Protocol: A protocol that transports IMG metadata from an IMG sender to IMG receiver(s) [4].

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) [4].

IMG Transfer: A transfer of IMG metadata from an IMG sender to IMG receiver(s).

3. IMG Common Framework Model

Two common elements are found in all existing IMG candidate cases: the need to describe the services and the need to deliver the descriptions. In some cases, the descriptions provide multicast addresses and thus are part of the transport configuration. In other cases, descriptions are specific to the media application and may be meant for either human or machine consumption. Thus, the technologies can be roughly divided into three areas:

- o Application-specific Metadata: data describing the content of services and media which are both specific to certain applications and generally human readable.
- o Delivery Descriptions: the descriptions (metadata) that are essential to enable (e.g., multicast) delivery. These include framing (headers) for application-specific metadata, the metadata element identification and structure, and fundamental session data.
- o Delivery Protocols: the methods and protocols to exchange descriptions between the senders and the receivers. An IMG transport protocol consists of two functions: carrying IMG metadata from an IMG sender to an IMG receiver and controlling an IMG transport protocol. These functions are not always exclusive; therefore, some messages may combine control messages and metadata carriage functions to reduce the amount of the messaging.

3.1. IMG Data Types

A data model is needed to precisely define the terminology and relationships between sets, supersets, and subsets of metadata. A precise data model is essential for the implementation of IMGs, although it is not within the scope of this framework and requires a separate specification. However, there are three IMG data types in general: Complete IMG Descriptions, Delta IMG Descriptions, and IMG Pointers.

3.1.1. Complete IMG Description

A Complete IMG Description provides a self-contained set of metadata for one media object or service, i.e., it does not need additional information from any other IMG element. This is not to be confused with "complete IMG metadata", which, although vaguely defined here,

represents the complete IMG metadata database of an IMG sender (or related group of IMG senders -- potentially the complete Internet IMG knowledge). An IMG sender will generally deliver only subsets of metadata from its complete database in a particular IMG transport session.

3.1.2. Delta IMG Description

A Delta IMG Description provides only part of a Complete IMG Description, defining the difference from a previous version of the Complete IMG Description. Delta IMG Descriptions may be used to reduce network resource usage, for instance, when data consistency is essential and small and frequent changes occur to IMG elements. Thus, this description does not represent a complete set of metadata until it is combined with other metadata that may already exist or arrive in the future.

3.1.3. IMG Pointer

An IMG Pointer identifies or locates metadata. This may be used to separately obtain metadata (Complete or Delta IMG Descriptions) or perform another IMG management function such as data expiry (and erasure). The IMG Pointer may be used to reference IMG metadata elements within the IMG transport session and across IMG transport sessions. This pointer type does not include IMG metadata per se (although it may also appear as a data field in Complete or Delta IMG Descriptions).

3.2. IMG Entities

There are several fundamental IMG entities that indicate the capability to perform certain roles. An Internet host involved in IMG operations may adopt one or more of these roles, which are defined in more detail in [Section 3.3](#).

IMG Announcer: sends IMG ANNOUNCE

IMG Listener: receives IMG ANNOUNCE

IMG Querier: sends IMG QUERY and receives IMG RESOLVE

IMG Resolver: receives IMG QUERY then sends IMG RESOLVE

IMG Subscriber: sends IMG SUBSCRIBE then receives IMG NOTIFY

IMG Notifier: receives IMG SUBSCRIBE then sends IMG NOTIFY

Figure 1 shows the relationship between IMG entities and the IMG sender and receiver.

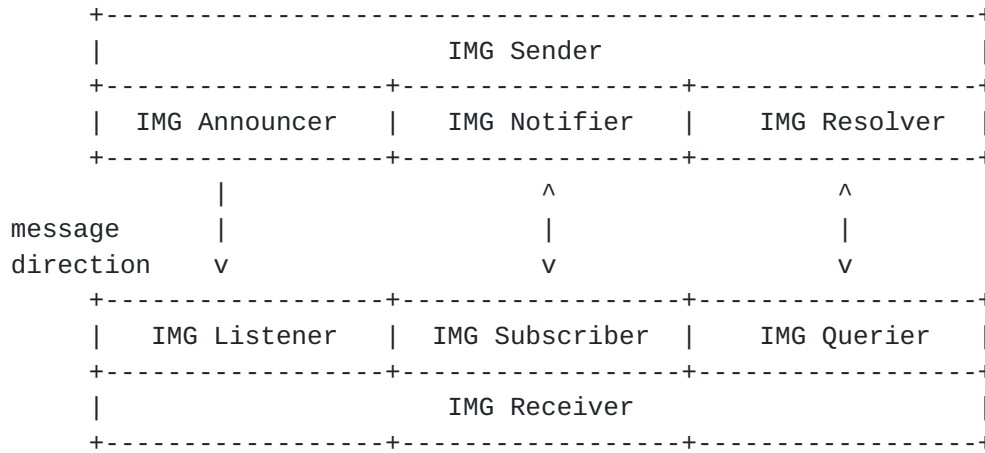


Figure 1: Relationship between IMG Entities, Senders, and Receivers

3.3. Operation Set for IMG Delivery

A finite set of operations both meets the IMG requirements [4] and fits the roles of existing protocols. These are crystallized in the next few sections.

3.3.1. IMG ANNOUNCE

When an IMG receiver participates in unidirectional communications (e.g., over satellite, terrestrial radio, and wired multicast networks), an IMG receiver may not need to send any IMG message to an IMG sender prior to IMG metadata delivery. In this case, an IMG sender can initiate unsolicited distribution for IMG metadata and an IMG sender is the only entity that can maintain the distribution (this includes scenarios with multiple and cooperative IMG senders). This operation is useful when there are large numbers of IMG receivers or the IMG receivers do not have a guaranteed uplink connection to the IMG sender. The IMG sender may also include authentication data in the ANNOUNCE operation so that IMG receivers may check the authenticity of the metadata. This operation can carry any of the IMG data types.

There is no restriction to prevent IMG ANNOUNCE from being used in an asynchronous solicited manner, where a separate operation (possibly out of band) enables IMG receivers to subscribe/register to the IMG ANNOUNCE operation.

3.3.2. IMG QUERY

If an IMG receiver needs to obtain IMG metadata, an IMG receiver can use an IMG QUERY operation and initiate a receiver-driven IMG transport session. The IMG receiver expects a synchronous response to the subsequent request from the IMG sender. This operation can be used where a bidirectional transport network is available between the IMG sender and receiver. Some IMG receivers may want to obtain IMG metadata when network connectivity is available or just to avoid caching unsolicited IMG metadata. The IMG receiver must indicate the extent and data type of metadata wanted in some message in the operation. The extent indicates the number and grouping of metadata descriptions. In some cases, it may be feasible to request an IMG sender's complete metadata collection.

3.3.3. IMG RESOLVE

An IMG sender synchronously responds, and sends IMG metadata, to an IMG QUERY that has been sent by an IMG receiver. This operation can be used where a bidirectional transport network is available between the IMG sender and receiver. If the IMG QUERY specifies a subset of IMG metadata (extent and data type) that is available to the IMG sender, the IMG sender can resolve the query; otherwise, it should indicate that it is not able to resolve the query. The IMG sender may authenticate the IMG receiver during the QUERY operation to determine if the IMG receiver is authorized to receive that metadata. The sender may also include authentication data in the RESOLVE operation so that IMG receivers may check the authenticity of the metadata. This operation may carry any of the IMG data types.

3.3.4. IMG SUBSCRIBE

If an IMG receiver wants to be notified when the IMG metadata it holds becomes stale, the IMG receiver can use the IMG SUBSCRIBE operation in advance in order to solicit IMG NOTIFY messages from an IMG sender.

This operation may provide the IMG sender with specific details of which metadata or notification services it is interested in the case where the IMG sender offers more than the simplest "all data" service. This operation implicitly provides the functionality of unsubscribing to inform an IMG sender that an IMG receiver wishes to stop getting certain (or all) notifications. It should be noted that unsubscription may be provided implicitly by the expiry (timeout) of a subscription before it is renewed.

Since the IMG receiver does not know when metadata will be updated and the notify message will arrive, this operation does not synchronize with the notify messages. The IMG receiver may wait for notify messages for a long time. The IMG sender may authenticate the IMG receiver to check whether an IMG SUBSCRIBE operation is from an authorized IMG receiver.

3.3.5. IMG NOTIFY

An IMG NOTIFY is used asynchronously in response to an earlier IMG SUBSCRIBE. An IMG NOTIFY operation indicates that updated IMG metadata is available or part of the existing IMG metadata is stale. An IMG NOTIFY may be delivered more than once during the time an IMG SUBSCRIBE is active. This operation may carry any of the IMG data types. The IMG sender may also include authentication data in the IMG NOTIFY operation so that IMG receivers may check the authenticity of the messages.

3.3.6. Binding between IMG Operations and Data Types

There is a need to provide a binding between the various IMG operations and IMG data types to allow management of each discrete set of IMG metadata transferred using an IMG operation. This binding must be independent of any particular metadata syntax used to represent a set of IMG metadata, as well as be compatible with any IMG transport protocol. The binding must uniquely identify the set of IMG metadata delivered within an IMG transfer, regardless of the metadata syntax used. The uniqueness may only be needed within the domains the metadata is used, but this must enable globally unique identification to support Internet usage. Care should be taken that scope- and domain-specific identifiers do not leak outside the scope; using globally unique identifiers such as URLs avoids these problems.

The binding must provide versioning to the transferred IMG metadata so that changes can be easily handled and stale data identified, and give temporal validity of the transferred IMG metadata. It must invalidate the IMG metadata by indicating an expiry time, and may optionally provide a time (presumably in the future) from when the IMG metadata becomes valid. The temporal validity of a metadata object may need to be updated later. Furthermore, the binding must be independent of any specific metadata syntax used for the IMG metadata, in the sense that no useful syntax should be excluded.

3.4. Overview of Protocol Operations

Figure 2 gives an overview of the relationship between transport cases, IMG operations, and IMG data types. It is not a protocol stack. Generalized multicast point-to-multipoint (P-to-M) and unicast point-to-point (P-to-P) transports are shown.

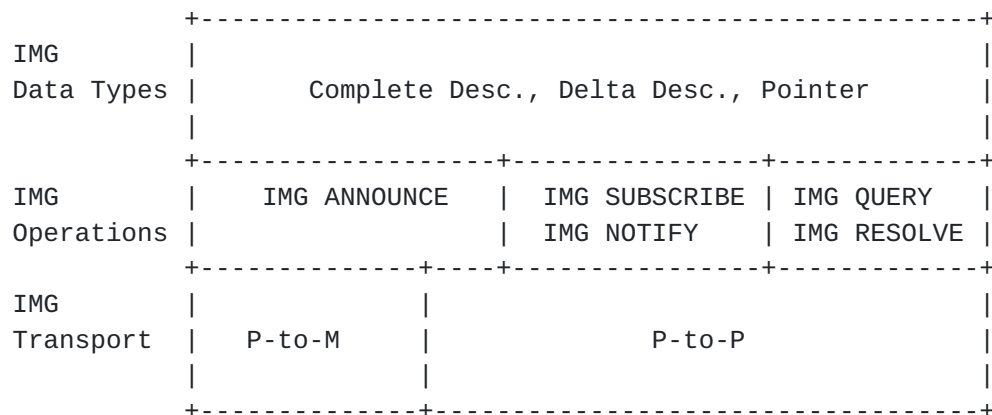


Figure 2: IMG Transport, IMG Operations, and IMG Data Types

4. Deployment Scenarios for IMG Entities

This section provides some basic deployment scenarios for IMG entities that illustrate common threads from protocols to use cases. For the purposes of clarity, this document presents the simple dataflow from an IMG sender to an IMG receiver, as shown in Figure 3.




Figure 3: A Simple IMG Sender to IMG Receiver Relationship

4.1. Intermediary Cases

Some IMG metadata may be distributed to a large number of IMG receivers, for example, when public metadata is distributed to all IMG receivers of a certain group. This kind of IMG metadata may be distributed from one IMG sender to multiple IMG receivers (Figure 4) or may be relayed across a set of IMG transceivers that receive the IMG metadata, possibly filter or modify its content, and then forward it.

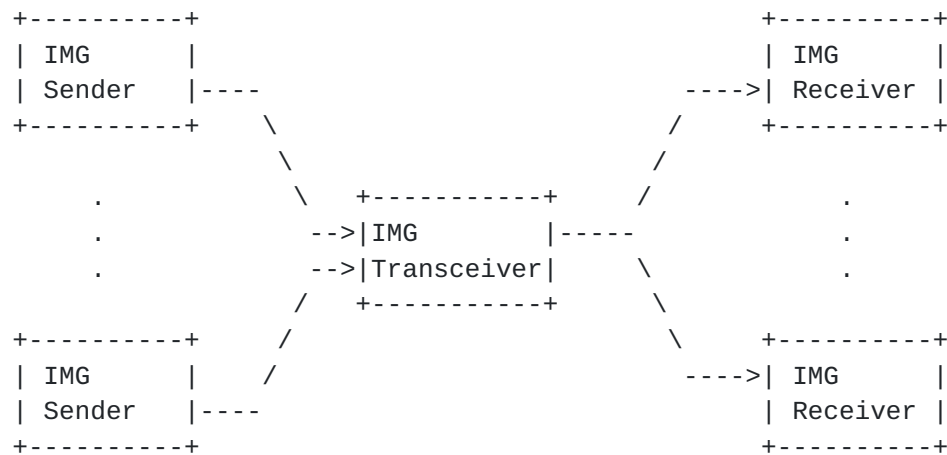


Figure 4: A Relay Network with an IMG Transceiver

IMG senders and receivers are logical functions, and it is possible for some or all hosts in a system to perform both roles, as, for instance, in many-to-many communications or where a transceiver is used to combine or aggregate IMG metadata for some IMG receivers. An IMG receiver may be allowed to receive IMG metadata from any number of IMG senders.

IMG metadata is used to find, obtain, manage, and play media content. IMG metadata may be modified during IMG transfer. For example, a server may use IMGs to retrieve media content via unicast and then make it available at scheduled times via multicast, thus requiring a change of the corresponding metadata. IMG transceivers may add or delete information or aggregate IMG metadata from different IMG senders. For example, a rating service may add its own content ratings or recommendations to existing metadata. An implication of changing (or aggregating) IMG metadata from one or more IMG senders is that the original authenticity is lost. Thus, it may be beneficial to sign fragments so that the intermediary can replace a fragment without changing the authenticity of the remainder. For example, smaller fragments may be appropriate for more volatile parts, and larger ones may be appropriate for stable parts.

4.2. One-to-Many Unidirectional Multicast

The one-to-many unidirectional multicast case implies many IMG receivers and one or more IMG senders implementing IMG announcer and IMG listener operations as shown in Figure 5.

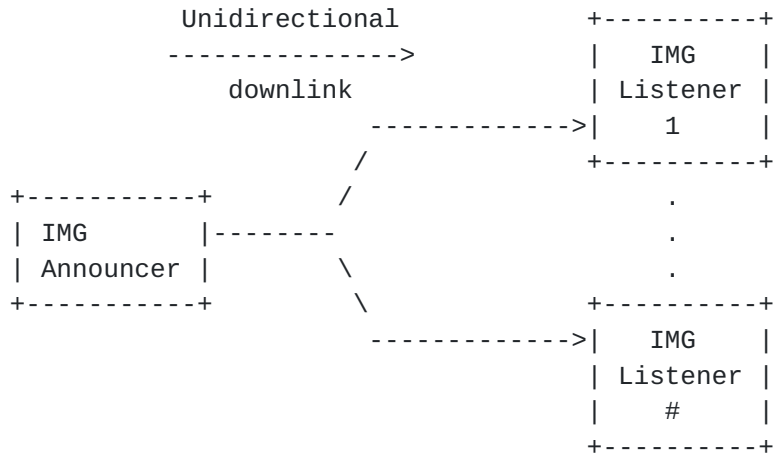


Figure 5: IMG Unidirectional Multicast Distribution Example

Note, as defined in the IMG requirement REL-4 [4], an IMG transport protocol MUST support reliable message exchange. This includes the one-to-many unidirectional multicast case; however, the mechanism to provide this is beyond the scope of this document.

4.3. One-to-One Bidirectional Unicast

In the one-to-one bidirectional unicast case, both query/resolve (Figure 6) and subscribe/notify (Figure 7) message exchange operations are feasible.




Figure 6: Query/Resolve Sequence Example

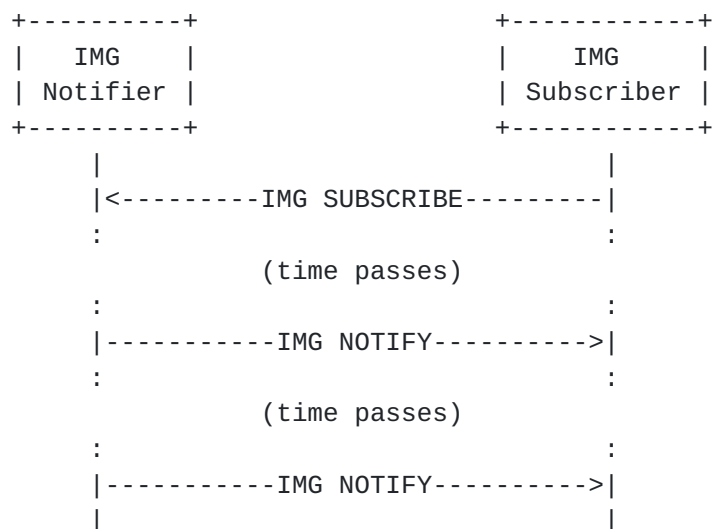


Figure 7: Subscribe/Notify Sequence Example

4.4. Combined Operations with Common Metadata

As shown in Figure 8, a common data model for multiple protocol operations allows a diverse range of IMG senders and receivers to provide consistent and interoperable sets of IMG metadata.

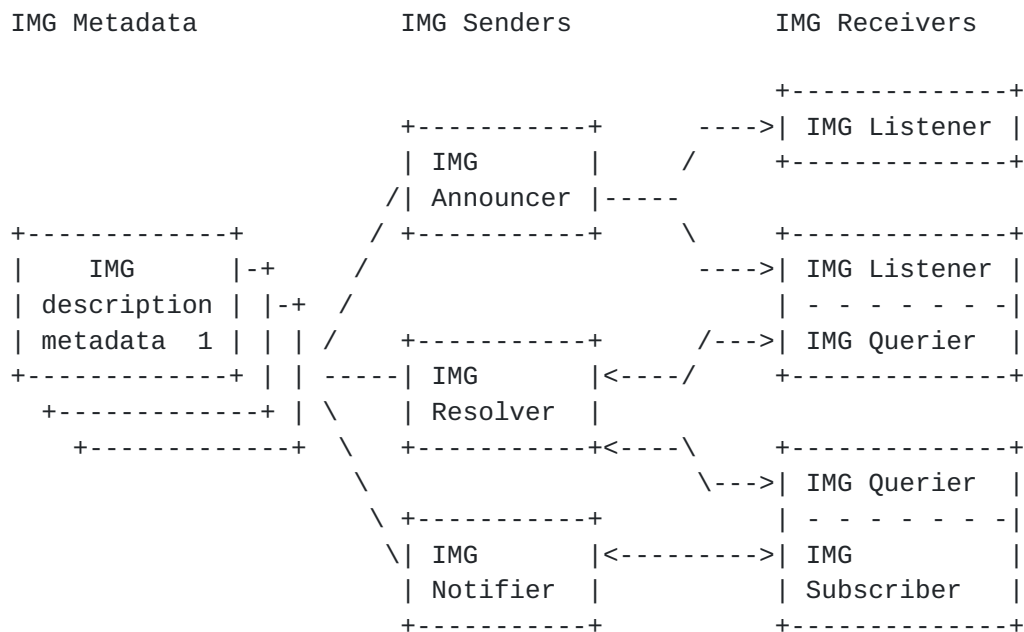


Figure 8: Combined System with Common Metadata

5. Applicability of Existing Protocols to the Proposed Framework Model

5.1. Existing Standards Fitting the IMG Framework Model

SDP: The SDP format [2] could be used to describe session-level parameters (e.g., scheduling, addressing, and the use of media codecs) to be included in Complete IMG Descriptions. Although there are extension points in SDP allowing the format to be extended, there are limitations in the flexibility of this extension mechanism. However, SDP syntax cannot provide IMG Descriptions and IMG Pointers without significant overhead. It is expected that the information conveyed by SDP is just a small subset of IMG metadata; thus, the use of SDP for other than session parameters may not be reasonable.

SDPng [3]: Similar to SDP, this format could also be used for representing session-level parameters of IMG metadata. Compared to SDP, the XML-based format of SDPng should be much more flexible and allow extensions and integration with other description formats.

MPEG-7: Descriptions based on the MPEG-7 standard [5] could provide application-specific metadata describing the properties of multimedia content beyond parameters carried in SDP or SDPng descriptions. MPEG-7 provides a machine-readable format of representing content categories and attributes, helping end-users or receiving software in choosing content for reception. MPEG-7 is based on XML, so it is well suited to be combined with other XML-based formats such as SDPng.

TV-Anytime: The TV-Anytime Forum [6] provides descriptions based on XML schema for TV-specific program guides. TV-Anytime uses the MPEG-7 User description profile to a limited extent, only for user preferences and usage history, and also a TV-Anytime-specific data model for other schema. These are optimized to describe broadcast schedules, on-demand program guides and program events.

HTTP: The HTTP protocol [7] can be used as a bidirectional unicast IMG transport protocol. Being a request-reply-oriented protocol, HTTP is well suited for implementing synchronous operations such as QUERY, RESOLVE, and even SUBSCRIBE. However, HTTP does not provide asynchronous operations such as ANNOUNCE and NOTIFY and to implement asynchronous operations using HTTP, IMG receivers should poll the IMG sender periodically. Thus, by itself, HTTP is not sufficient to fulfill all of the IMG requirements [4] in a unicast deployment.

Session Announcement Protocol (SAP): The announcement mechanism provided by SAP [8] provides unidirectional delivery of session discovery information. Although SDP is the default payload format of SAP, the use of a MIME type identifier for the payload allows

arbitrary payload formats to be used in SAP messages. Thus, SAP could be used to implement the multicast and unicast IMG ANNOUNCE and IMG NOTIFY operations.

However, SAP lacks scalable and efficient reliability, extensibility for payload size, and congestion control, and only one description is allowed per SAP message due to lack of payload segmentation.

In principle, SAP could be extended to get around its limitations. However, the amount of changes needed in SAP to address all of the above limitations would effectively result in a new protocol. Due to these limitations, the use of SAP as an IMG transport protocol is not recommended.

SIP: The SIP-specific event mechanism described in [RFC 3265](#) [9] provides a way to implement IMG SUBSCRIBE and IMG NOTIFY operations via a bidirectional unicast connection. However, there are scalability problems with this approach, as [RFC 3265](#) currently does not consider multicast.

Real Time Streaming Protocol (RTSP): The RTSP protocol [10] defines a retrieval-and-update notification mechanism, named DESCRIBE and ANNOUNCE, for the description of a presentation or media object in order to initialize a streaming session. These methods are a subset of the entire streaming control operations in RTSP; thus, these could not be available for individual mechanisms. However, the DESCRIBE method in RTSP could be used to instantiate IMG QUERY, IMG RESOLVE, and IMG SUBSCRIBE, and the RTSP ANNOUNCE could be used to instantiate an IMG NOTIFY for a streaming session controlled by RTSP.

5.2. IMG Mechanism Needs Which Are Not Met by Existing Standards

Several needs result from the IMG requirements, framework model, and existing relevant mechanisms as already shown in this document. Four specific groupings of work are readily apparent: (a) specification of an adequate multicast- and unidirectional-capable announcement protocol; (b) specification of the use of existing unicast protocols to enable unicast subscribe and announcement/notification functionality; (c) specification of the metadata envelope that is common to, and independent of, the application metadata syntax(es) used; and (d) agreement on basic metadata models to enable interoperability testing of the above. The following sections describe each of these.

5.2.1. A Multicast Transport Protocol

SAP is currently the only open standard protocol suited to the unidirectional/multicast delivery of IMG metadata. As discussed, it fails to meet the IMG requirements in many ways and, since it is not designed to be extensible, we recognize that a new multicast transport protocol for announcements needs to be specified to meet IMG needs. This protocol will be essential to IMG delivery for unidirectional and multicast deployments.

The Asynchronous Layered Coding (ALC) [11] protocol from the IETF Reliable Multicast Transport (RMT) working group is very interesting as it fulfills many of the requirements, is extensible, and has the ability to 'plug-in' both FEC (Forward Error Correction, for reliability) and CC (Congestion Control) functional blocks. It is specifically designed for unidirectional multicast object transport. ALC is not fully specified, although the RMT working group had a fully specified protocol using ALC called FLUTE (File Delivery over Unidirectional Transport) [12]. FLUTE seems to be the only fully specified transport and open specification on which a new IMG announcement protocol could be designed. Thus, we recommend that ALC and FLUTE be the starting points for this protocol's design.

Developing a new protocol from scratch, or attempting to improve SAP, is also feasible, although it would involve repeating many of the design processes and decisions already made by the IETF for ALC. In particular, any announcement protocol must feature sufficient scalability, flexibility, and reliability to meet IMG needs. Also, the IMG ANNOUNCE operation must be supported and IMG NOTIFY capability could be investigated for both hybrid unicast-multicast and unidirectional unicast systems.

5.2.2. Usage of Unicast Transport Protocols

A thorough description of the use of existing unicast protocols is essential for the use of IMGs in a unicast point-to-point environment. Such a specification has not been published, although several usable unicast transport protocols and specifications can be harnessed for this (SIP [13], SIP events [9], HTTP [7], etc.). In particular, both IMG SUBSCRIBE-NOTIFY and IMG QUERY-RESOLVE operation pairs must be enabled. We anticipate that the IMG QUERY-RESOLVE operation can be achieved using HTTP, although other transport protocol options may be beneficial to consider too.

5.2.3. IMG Envelope

An IMG envelope provides the binding between IMG operations and data types. Such a binding can be realized by defining a common minimal set of information needed to manage IMG metadata transfers, and by including this information with any set of IMG metadata delivered to IMG receivers.

Four options for IMG metadata transfer envelope delivery are feasible:

1. Embedding in a transport protocol header. This can be done with either header extensions of existing protocols, or newly defined header fields of a new (or new version of a) transport protocol. However, multiple methods for the variety of transport protocols would hinder interoperability and transport protocol independence.
2. A separate envelope object, which points to the IMG metadata 'object', delivered in-band with the metadata transport protocol session. This might complicate delivery as the envelope and 'service' metadata objects would have to be bound, e.g., by pairing some kind of transport object numbers (analogous to port number pairs sometimes used for RTP and RTCP [14]). This would also enable schemes that deliver envelope and metadata 'objects' by different media, also using more than a single transport protocol.
3. A metadata wrapper that points to and/or embeds the service metadata into its 'super-syntax'. For example, XML would enable embedding generic text objects.
4. Embedding in the metadata itself. However, this requires a new field in many metadata syntaxes and would not be feasible if a useful syntax were not capable of extensibility in this way. It also introduces a larger 'implementation interpretation' variety, which would hinder interoperability. Thus, this option is not recommended.

It is likely that more than one of these options will fulfill the needs of IMGs, so the selection, and possibly optimization, is left for subsequent specification and feedback from implementation experience. Such a specification is essential for IMG delivery.

When there are superset/subset relations between IMG Descriptions, it is assumed that the IMG Descriptions of the subset inherit the parameters of the superset. Thus, an IMG metadata transfer envelope carrying the IMG Descriptions of a superset may implicitly define

parameters of IMG Descriptions belonging to its subset. The relations between IMG Descriptions may span from one envelope to another according to a data model definition.

5.2.4. Metadata Data Model

A structured data model would allow reuse and extension of the set of metadata and may enable use of multiple syntaxes (SDP, MPEG-7, etc.) as part of the same body of IMG metadata.

For the successful deployment of IMGs in various environments, further work may be needed to define metadata and data models for application-specific requirements. Existing (and future) work on these would need to be mapped to the IMG data types and use of the IMG transfer envelope concept as described above.

This document is a framework for the delivery of IMG metadata and thus further discussion on the definition data models for IMGs is beyond its scope.

6. Security Considerations

The IMG framework is developed from the IMG requirements document [4], and so the selection of specific protocols and mechanism for use with the IMG framework must also take into account the security considerations of the IMG requirements document. This framework document does not mandate the use of specific protocols. However, an IMG specification would inherit the security considerations of specific protocols used.

Protocol instantiations that are used to provide IMG operations will have very different security considerations depending on their scope and purpose. However, there are several general issues that are valuable to consider and, in some cases, provide technical solutions for. These are described below.

Individual and group privacy: Customized IMG metadata may reveal information about the habits and preferences of a user and may thus deserve confidentiality protection, even if the original information were public. Protecting this metadata against snooping requires the same actions and measures as for other point-to-point and multicast Internet communications. Naturally, the risk of snooping depends on the amount of individual or group personalization the IMG metadata contains.

IMG authenticity: In some cases, the IMG receiver needs to be assured of the sender or origin of IMG metadata or its modification history. This can prevent denial-of-service or hijacking attempts that give an

IMG receiver incorrect information in or about the metadata, thus preventing successful access of the media or directing the IMG receiver to the incorrect media possibly with tasteless material.

IMG receiver authorization: Some or all of any IMG sender's metadata may be private or valuable enough to allow access to only certain IMG receivers and thus make it worth authenticating users. Encrypting the data is also a reasonable step, especially where group communications methods results in unavoidable snooping opportunities for unauthorized nodes.

Unidirectional specifics: A difficulty that is faced by unidirectional delivery operations is that many protocols providing application-level security are based on bidirectional communications. The application of these security protocols in case of strictly unidirectional links is not considered in the present document.

Malicious code: Currently, IMGs are not envisaged to deliver executable code at any stage. However, as some IMG transport protocols may be capable of delivering arbitrary files, it is RECOMMENDED that the IMG operations do not have write access to the system or any other critical areas.

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] Kutscher, D., Ott, J., and C. Bormann, "Session description and capability negotiation", Work in Progress, October 2003.
- [4] Nomura, Y., Walsh, R., Luoma, J-P., Ott, J., and H. Schulzrinne, "Requirements for Internet Media Guides", Work in Progress, December 2005.
- [5] "Multimedia content description interface -- Part 1: Systems", ISO/IEC 15938-1, July 2002.
- [6] TV-Anytime Forum, "Broadcast and On-line Services: Search, select, and rightful use of content on personal storage systems ("TV-Anytime Phase 1"); Part 2: System description," ETSI-TS 102 822-2: System Description, V1.1.1, October 2003.

- [7] 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.
- [8] Handley, M., Perkins, C., and E. Whelan, "Session Announcement Protocol", [RFC 2974](#), October 2000.
- [9] Roach, A., "Session Initiation Protocol (SIP)-Specific Event Notification", [RFC 3265](#), June 2002.
- [10] Schulzrinne, H., Rao, A., and R. Lanphier, "Real Time Streaming Protocol (RTSP)", [RFC 2326](#), April 1998.
- [11] Luby, M., Gemmell, J., Vicisano, L., Rizzo, L., and J. Crowcroft, "Asynchronous Layered Coding (ALC) Protocol Instantiation", [RFC 3450](#), December 2002.
- [12] Paila, T., Luby, M., Lehtonen, R., Roca, V., and R. Walsh, "FLUTE - File Delivery over Unidirectional Transport", [RFC 3926](#), October 2004.
- [13] 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.
- [14] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, [RFC 3550](#), July 2003.

9. Acknowledgements

The authors would like to thank Spencer Dawkins, Jean-Pierre Evain, Ted Hardie, Petri Koskelainen, Joerg Ott, Colin Perkins, Toni Paila, and Magnus Westerlund for their excellent ideas and input to this document.

Authors' Addresses

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

E-Mail: nom@flab.fujitsu.co.jp

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

E-Mail: rod.walsh@nokia.com

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

E-Mail: juha-pekka.luoma@nokia.com

Hitoshi Asaeda
Keio University
Graduate School of Media and Governance
5322 Endo, Fujisawa, 252-8520 Kanagawa,
Japan

E-Mail: asaeda@wide.ad.jp

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

E-Mail: 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).

