

**Providing integrated services over low-bitrate links
draft-ietf-issll-isslow-01.txt**

Status of this memo

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as ``work in progress.''

To learn the current status of any Internet-Draft, please check the ``[l1d-abstracts.txt](#)'' listing contained in the Internet-Drafts Shadow Directories on [ftp.is.co.za](#) (Africa), [nic.nordu.net](#) (Europe), [munnari.oz.au](#) (Pacific Rim), [ds.internic.net](#) (US East Coast), or [ftp.isi.edu](#) (US West Coast).

Distribution of this document is unlimited.

Abstract

This document describes an architecture for providing integrated services over low-bitrate links, such as modem lines, ISDN B-channels, and sub-T1 links. It covers only the lower parts of the Internet Multimedia Conferencing Architecture [\[1\]](#); additional components required for application services such as Internet Telephony (e.g., a session initiation protocol) are outside the scope of this document. The main components of the architecture are: a real-time encapsulation format for asynchronous and synchronous low-bitrate links, a header compression architecture optimized for real-time flows, elements of negotiation protocols used between routers (or between hosts and routers), and announcement protocols used by applications to allow this negotiation to take place.

This document is a product of the IETF ISSLL working group. Comments are solicited and should be addressed to the working group's mailing list at issll@mercury.lcs.mit.edu and/or the author.

1. Introduction

As an extension to the ``best-effort'' services the Internet is well-known for, additional types of services (``integrated services'') that support the transport of real-time multimedia information are being developed for and deployed in the Internet. Important elements of this development are:

- parameters for forwarding mechanisms that are appropriate for real-time information (in the intserv working group of the IETF)
- a setup protocol that allows establishing special forwarding treatment for real-time information flows (RSVP [\[4\]](#), in the rsvp working group of the IETF)
- a transport protocol for real-time information (RTP/RTCP [\[6\]](#), in the avt working group of the IETF).

In addition to these elements at the network and transport levels of the Internet Multimedia Conferencing Architecture [\[1\]](#), further components are required to define application services such as Internet Telephony, e.g., protocols for session initiation and control. These components are outside the scope of this document.

Up to now, the newly developed services could not (or only very inefficiently) be used over forwarding paths that include low-bitrate links such as 14.4 and 28.8 kbit/s modems, 56 and 64 kbit/s ISDN B-channels, or even sub-T1 links. The encapsulation formats used on these links are not appropriate for the simultaneous transport of arbitrary data and real-time information that has to meet stringent delay requirements. A 1500 byte packet on a 28.8 kbit/s modem link makes this link unavailable for the transmission of real-time information for about 400 ms. This adds a worst-case delay that causes real-time applications to operate with round-trip delays on the order of at least a second -- unacceptable for real-time conversation. In addition, the header overhead associated with the protocol stacks used is prohibitive on low-bitrate links, where compression down to a few dozen bytes per real-time information packet is often desirable. E.g., the overhead of at least 44 (4+20+8+12) bytes for HDLC/PPP, IP, UDP, and RTP completely overshadows the 19.75 bytes needed for a G.723.1 ACELP audio frame -- a 14.4 kbit/s link is completely consumed by this header overhead alone at 40 real-time frames per second (i.e., at 25 ms packetization delay for one stream or 50 ms for two streams, with no space left for data, yet). While the header overhead can be reduced by combining several real-time information frames into one packet, this increases the delay incurred while filling that packet and further detracts from the goal of real-time transfer of multi-media information over the Internet.

This document describes an approach for addressing these problems.
The main components of the architecture are:

- a real-time encapsulation format for asynchronous and synchronous low-bitrate links,
- a header compression architecture optimized for real-time flows,
- elements of negotiation protocols used between routers (or between hosts and routers), and
- announcement protocols used by applications to allow this negotiation to take place.

2. Design Considerations

The main design goal for an architecture that addresses real-time multimedia flows over low-bitrate links is that of minimizing the end-to-end delay. More specifically, the worst case delay (after removing possible outliers, which are equivalent to packet losses from an application point of view) is what determines the playout points selected by the applications and thus the delay actually perceived by the user.

In addition, any such architecture should obviously undertake every attempt to maximize the bandwidth actually available to media data; overheads must be minimized.

An important component of the integrated services architecture is the provision of reservations for real-time flows. One of the problems that systems on low-bitrate links (routers or hosts) face when performing admission control for such reservations is that they must translate the bandwidth requested in the reservation to the one actually consumed on the link. Methods such as data compression and/or header compression can reduce the requirements on the link, but admission control can only make use of the reduced requirements in its calculations if it has enough information about the data stream to know how effective the compression will be. One goal of the architecture therefore is to provide the integrated services admission control with this information. A beneficial side effect may be to allow the systems to perform better compression than would be possible without this information. This may make it worthwhile to provide this information even when it is not intended to make a reservation for a real-time flow.

3. The Need for a Concerted Approach

Many technical approaches come to mind for addressing these problems, in particular a new form of low-delay encapsulation to address delay

and header compression methods to address overhead. This section shows that any single one of these techniques is not sufficient to solve the problem.

3.1. Real-Time Encapsulation

The purpose of defining a real-time link-layer encapsulation protocol is to be able to introduce newly arrived real-time packets in the link-layer data stream without having to wait for the currently transmitted (possibly large) packet to end. Obviously, a real-time encapsulation must be part of any complete solution as the problem of delays induced by large frames on the link can only be solved on this layer.

To be able to switch to a real-time packet quickly in an interface driver, it is first necessary to identify packets that belong to real-time flows. This can be done using a heuristic approach (e.g., favor the transmission of highly periodic flows of small packets transported in IP/UDP). Preferably, one also could make use of the protocol defined for identifying flows that require special treatment, i.e. RSVP. Of the two service types defined for use with RSVP now, the guaranteed service will only be available in certain environments; for this and various other reasons, the service type chosen for many adaptive audio/video applications will be the controlled-load service. Controlled-load does not provide control parameters for target delay; this makes it very hard to identify those packet streams that would benefit most from being transported in a real-time encapsulation format. This calls for a way to provide additional parameters in integrated services flow setup protocols to control the real-time encapsulation.

Real-time encapsulation is not sufficient on its own, however: Even if the relevant flows can be appropriately identified for real-time treatment, most applications simply are not possible on low-bitrate links with the header overhead implied by the combination of HDLC/PPP, IP, UDP, and RTP, i.e. they absolutely require header compression.

3.2. Header Compression

Header compression can be performed in a variety of elements and at a variety of levels in the protocol architecture. As internet telephony vendors ship applications, the approach that is most obvious to them is to reduce overhead by performing header compression at the application level, i.e. above transport protocols such as UDP[1].

Generally, header compression operates by installing state at both ends of a path that allows the receiving end to reconstruct information omitted at the sending end. Many good techniques for header compression ([RFC 1144](#), [2]) operate on the assumption that the path will not reorder the frames generated. This assumption does not

[[1](#)] or actually by using a non-standard, efficiently coded header in the first place.

hold for end-to-end compression; therefore additional overhead is required for resequencing state changes and compressed packets making use of these state changes.

Assume that a very good application level header compression solution for RTP flows could be able to save 11 out of the 12 bytes of an RTP header [3]. Even this perfect solution only reduces the total header overhead by 1/4. It would have to be deployed in all applications, even those that operate on systems that are attached to higher-bitrate links.

Because of this limited effectiveness, the AVT group that is responsible for RTP within the IETF has decided to not further pursue application level header compression.

For router and IP stack vendors, the obvious approach is to define header compression that can be negotiated between peer routers.

Advanced header compression techniques now being defined in the IETF [2] certainly can relieve the link from significant parts of the IP/UDP overhead (i.e., of 28 of the 44 bytes mentioned above).

One of the design principles of the header compression contemplated for IPv6 is that it stops at layers the semantics of which cannot be inferred from information in lower layer (outer) headers. Therefore, the original IPv6 header compression cannot compress the data that is contained within UDP packets.

Any additional header compression technique runs into a problem: If it assumes specific application semantics (i.e., those of RTP and a payload data format) based on heuristics, it runs the risk of being triggered falsely and (e.g. in case of packet loss) reconstructing packets that are catastrophically incorrect for the application actually being used. A header compression technique that can be operated based on heuristics but does not cause incorrect decompression even if the heuristics failed is described in [7].

With all of these techniques, the total IP/UDP/RTP header overhead for an audio stream can be reduced to two bytes per packet. This technology need only be deployed at bottleneck links; high-speed links can transfer the real-time streams without routers or switches expending CPU cycles to perform header compression.

4. Principles of Real-Time Encapsulation for Low-Bitrate Links

The main design goal for a real-time encapsulation is to minimize the delay incurred by real-time packets that become available for sending while a long data packet is being sent. To achieve this, the

encapsulation must be able to either abort or suspend the transfer of the long data packet. An additional goal is to minimize the overhead required for the transmission of packets from periodic flows. This

strongly argues for being able to suspend a packet, i.e. segment it into parts between which the real-time packets can be transferred.

Cell-oriented multiplexing techniques such as ATM that introduce regular points where cells from a different packet can be interpolated are too inefficient for low-bitrate links; also, they are not supported by chips used to support the link layer in low-bitrate routers and host interfaces.

Instead, the real-time encapsulation should as far as possible make use of the capabilities of the chips that have been deployed. On synchronous lines, these chips support HDLC framing; on asynchronous lines, an asynchronous variant of HDLC that usually is implemented in software is being used. Both variants of HDLC provide a delimiting mechanism to indicate the end of a frame over the link. The obvious solution to the segmentation problem is to combine this mechanism with an indication of whether the delimiter terminates or suspends the current packet.

This indication could be in an octet appended to each frame information field; however, seven out of eight bits of the octet would be wasted. Instead, the bit could be carried at the start of the next frame in conjunction with multiplexing information (PPP protocol identifier etc.) that will be required here anyway. Since the real-time flows will in general be periodic, this multiplexing information could convey (part of) the compressed form of the header for the packet. If packets from the real-time flow generally are of constant length (or have a defined maximum length that is often used), the continuation of the suspended packet could be immediately attached to it, without expending a further frame delimiter, i.e., the interpolation of the real-time packet would then have zero overhead. Since packets from low-delay real-time flows generally will not require the ability to be further suspended, the continuation bit could be reserved for the non-real-time packet stream.

One real-time encapsulation format with these (and other) functions is described in ITU-T H.223, the multiplex used by the H.324 videophone standard. It was investigated whether compatibility could be achieved with this specification, which will be used in future videophone-enabled (H.324 capable) modems. However, since the multiplexing capabilities of H.223 are limited to 15 schedules (definitions of sequences of packet types that can be identified in a multiplex header), for general Internet usage a superset or a more general encapsulation would have been required. Also, a PPP-style negotiation protocol was needed instead of using (and necessarily extending) ITU-T H.245 for setting the parameters of the multiplex. In the PPP context, the interactions with the encapsulations for data compression and link layer encryption needed to be defined (including

operation in the presence of padding). But most important, H.223 requires synchronous HDLC chips that can be configured to send frames without an attached CRC, which is not possible with all chips deployed in commercially available routers; so complete compatibility

was unachievable.

Instead of adopting H.223, it was decided to pursue an approach that is oriented towards compatibility both with existing hardware and existing software (in particular PPP) implementations. The next subsection groups these implementations according to their capabilities.

4.1. Implementation models

This section introduces a number of terms for types of implementations that are likely to emerge. It is important to have these different implementation models in mind as there is no single approach that fits all models best.

4.1.1. Sender types

There are two fundamental approaches to real-time transmission on low-bitrate links:

Sender type 1

The PPP real-time framing implementation is able to control the transmission of each byte being transmitted with some known, bounded delay (e.g., due to FIFOs). For example, this is generally true of PC host implementations, which directly access serial interface chips byte by byte or by filling a very small FIFO. For type 1 senders, a suspend/resume type approach will be typically used: When a long frame is to be sent, the attempt is to send it undivided; only if higher priority packets come up during the transmission will the lower-priority long frame be suspended and later resumed. This approach allows the minimum variation in access delay for high-priority packets; also, fragmentation overhead is only incurred when actually needed.

Sender type 2

With type 2 senders, the interface between the PPP real-time framing implementation and the transmission hardware is not in terms of streams of bytes, but in terms of frames, e.g., in the form of multiple (prioritized) send queues directly supported by hardware. This is often true of router systems for synchronous links, in particular those that have to support a large number of low-bitrate links. As type 2 senders have no way to suspend a frame once it has been handed down for transmission, they typically will use a queues-of-fragments approach, where long packets are always split into units that are small enough to maintain the access delay goals for higher-priority traffic. There is a trade-off between the variation in access delay resulting from a large fragment size and the overhead that is

incurred for every long packet by choosing a small fragment size.

4.1.2. Receiver types

Although the actual work of formulating transmission streams for real-time applications is performed at the sender, the ability of the receiver to immediately make use of the information received depends on its characteristics:

Receiver type 1

Type 1 receivers have full control over the stream of bytes received within PPP frames, i.e., bytes received are available immediately to the PPP real-time framing implementation (with some known, bounded delay e.g. due to FIFOs etc.).

Receiver type 2

With type 2 receivers, the PPP real-time framing implementation only gets hold of a frame when it has been received completely, i.e., the final flag has been processed (typically by some HDLC chip that directly fills a memory buffer).

4.2. Conclusion

As a result of the diversity in capabilities of current implementations, there are now two specifications for real-time encapsulation: One, the multi-class extension to the PPP multi-link protocol, is providing the solution for the queues-of-fragments approach by extending the single-stream PPP multi-link protocol by multiple classes [8]. The other encapsulation, PPP in a real-time oriented HDLC-like framing, builds on this specification and extends it by a way to dynamically delimit multiple fragments within one HDLC frame [9], providing the solution for the suspend/resume type approach.

5. Principles of Header Compression for Real-Time Flows

A good baseline for a discussion about header compression is in the IPv6 header compression draft [2]. The techniques used there can reduce the 28 bytes of IPv4/UDP header to about 6 bytes (depending on the number of concurrent streams); with the remaining 4 bytes of HDLC/PPP overhead and 12 bytes for RTP the total header overhead can be about halved but still exceeds the size of a G.723.1 ACELP frame. Note that, in contrast to IPv6 header compression, the environment discussed here assumes the existence of a full-duplex PPP link and thus can rely on negotiation where IPv6 header compression requires repeated transmission of the same information. (The use of the architecture of the present document with link layer multicasting has not yet been examined.)

Additional design effort was required for RTP header compression.
Applying the concepts of IPv6 header compression, of the (at least)
12 bytes in an RTP header, 7 bytes (timestamp, sequence, and marker

bit) would qualify as RANDOM; DELTA encoding cannot generally be used without further information since the lower layer header does not unambiguously identify the semantics and there is no TCP checksum that can be relied on to detect incorrect decompression. Only a more semantics-oriented approach can provide better compression (just as [RFC 1144](#) can provide very good compression of TCP headers by making use of semantic knowledge of TCP and its checksumming method).

For RTP packets, differential encoding of the sequence number and timestamps is an efficient approach for certain cases of payload data formats. E.g., speech flows generally have sequence numbers and timestamp fields that increase by 1 and by the frame size in timestamp units, resp.; the CRTP (compressed RTP) specification makes use of this relationship by encoding these fields only when the second order difference is non-zero [7].

6. Announcement Protocols Used by Applications

As argued, the compressor can operate best if it can make use of information that clearly identifies real-time streams and provides information about the payload data format in use.

If these systems are routers, this consent must be installed as router state; if these systems are hosts, it must be known to their networking kernels. Sources of real-time information flows are already describing characteristics of these flows to their kernels and to the routers in the form of TSspecs in RSVP PATH messages [4]. Since these messages make use of the router alert option, they are seen by all routers on the path; path state about the packet stream is normally installed at each of these routers that implement RSVP. Additional RSVP objects could be defined that are included in PATH messages by those applications that desire good performance over low-bitrate links; these objects would be coded to be ignored by routers that are not interested in them (class number 11bbbbbb).

Note that the path state is available in the routers even when no reservation is made; this allows informed compression of best-effort traffic. It is not quite clear, though, how path state could be teared down quickly when a source ceases to transmit.

7. Elements of Hop-By-Hop Negotiation Protocols

The IPv6 header compression draft attempts to account for simplex and multicast links by providing information about the compressed streams only in the forward direction. E.g., a full IP/UDP header must be sent after F_MAX_TIME (currently 3 seconds), which is a negligible

total overhead (e.g. one full header every 150 G.723.1 packets), but must be considered carefully in scheduling the real-time transmissions. Both simplex and multicast links are not prevailing

in the low-bitrate environment (although multicast functionality may become more important with wireless systems); in this document, we therefore assume full-duplex capability.

As compression techniques will improve, a negotiation between the two peers on the link would provide the best flexibility in implementation complexity and potential for extensibility. The peer routers/hosts can decide which real-time packet streams are to be compressed, which header fields are not to be sent at all, which multiplexing information should be used on the link, and how the remaining header fields should be encoded. PPP, a well-tried suite of negotiation protocols, is already used on most of the low-bitrate links and seems to provide the obvious approach. Cooperation from PPP is also needed to negotiate the use of real-time encapsulations between systems that are not configured to automatically do so. Therefore, PPP options that can be negotiated at the link setup (LCP) phase are included in [8], and [9] (and need to be put into [7], FIX ME).

8. Security Considerations

Header compression protocols that make use of assumptions about application protocols need to be carefully analyzed whether it is possible to subvert other applications by maliciously or inadvertently enabling their use.

It is generally not possible to do significant hop-by-hop header compression on encrypted streams. With certain security policies, it may be possible to run an encrypted tunnel to a network access server that does header compression on the decapsulated packets and sends them over an encrypted link encapsulation; see also the short mention of interactions between real-time encapsulation and encryption in [section 4](#) above. If the security requirements permit, a special RTP payload data format that encrypts only the data may preferably be used.

9. Author's Address

Carsten Bormann
Universitaet Bremen FB3 TZI
Postfach 330440
D-28334 Bremen, GERMANY
cabo@tzi.uni-bremen.de
phone +49.421.218-7024
fax +49.421.218-7000

Acknowledgements

Much of the early discussion that led to this document was done with
Scott Petrack of IBM Israel and Cary Fitzgerald of Cisco Systems.

Bormann

[Page 10]

Steve Casner, Mikael Degermark, Steve Jackowski, Dave Oran, and the other members of the ISSLL subgroup on low bitrate links (ISSLOW) have helped making this architecture a reality.

10. References

- [1] Mark Handley, Jon Crowcroft, Carsten Bormann, ``The Internet Multimedia Conferencing Architecture'', Internet-Draft [draft-ietf-mmusic-confarch-00.txt](#), Work in Progress, February 1996.
- [2] M. Degermark, B. Nordgren, S. Pink, ``Header Compression for IPv6'', Internet-Draft [draft-degermark-ipv6-hc-02.txt](#), Work in Progress, November 1996.
- [3] Scott Petrack, Ed Ellessen, ``Framework for C/RTP: Compressed RTP Using Adaptive Differential Header Compression'', contribution to the mailing list rem-conf@es.net, February 1996.
- [4] R. Braden, Ed., L. Zhang, S. Berson, S. Herzog, S. Jamin, Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification, Internet-Draft [draft-ietf-rsvp-spec-14.txt](#), Work in Progress, November 1996.
- [5] K. Sklower, B. Lloyd, G. McGregor, D. Carr, T. Coradetti, ``The PPP Multilink Protocol (MP)'', [RFC 1990](#), August 1996 (obsoletes [RFC1717](#)).
- [6] H. Schulzrinne, S. Casner, R. Frederick & V. Jacobson, ``RTP: A Transport Protocol for Real-Time Applications'', [RFC 1889](#), January 1996.
- [7] S. Casner, V. Jacobson, ``Compressing IP/UDP/RTP Headers for Low-Speed Serial Links'', Internet-Draft [draft-ietf-avt-crtp-01.txt](#), Work in Progress, November 1996.
- [8] C. Bormann, ``The Multi-Class Extension to Multi-Link PPP'', internet Draft [draft-ietf-issll-isslow-mcml-01.txt](#), Work in Progress, March 1997.
- [9] C. Bormann, ``PPP in a real-time oriented HDLC-like framing, internet Draft [draft-ietf-issll-isslow-rtf-00.txt](#), Work in Progress, March 1997.

