Network Working Group
Internet Draft
<draft-ietf-avt-hc-over-mpls-protocol-02.txt>
Expiration Date: June 2006

Jerry Ash
AT&T
Andrew Malis
Tellabs

December, 2005

Protocol Extensions for Header Compression over MPLS

Status of this Memo

Copyright Notice

Abstract

VoIP typically uses the encapsulation voice/RTP/UDP/IP. When MPLS labels are added, this becomes voice/RTP/UDP/IP/MPLS-labels. For an MPLS VPN, the packet header is typically 48 bytes, while the voice payload is often no more than 30 bytes, for example. Header compression can significantly reduce the overhead through various compression mechanisms. MPLS is used to route header-compressed (HC) packets over an MPLS LSP without compression/decompression cycles at each router. Such an HC over MPLS capability increases the bandwidth efficiency as well as processing scalability of the maximum number of

simultaneous compressed flows that use HC at each router. MPLS pseudowires are used to transport the HC context and other control messages between the ingress and egress MPLS label switched router (LSR), and the pseudowires define one or more point-to-point instances corresponding to each HC session at the header decompressor. Standard HC methods (e.g., ECRTP, ROHC, etc.) are re-used to determine the context.

Table of Contents

1. Introduction

Voice over IP (VoIP) typically uses the encapsulation voice/RTP/UDP/IP. When MPLS labels [RFC3031] are added, this becomes voice/RTP/UDP/IP/MPLS-labels. MPLS VPNs (e.g., [RFC2547]) use label stacking, and in the simplest case of IPv4 the total packet header is at least 48 bytes, while the voice payload is often no more than 30 bytes, for example. When IPv6 is used, the relative size of the header in comparison to the payload is even greater. The interest in header compression is to exploit the possibility of significantly reducing the overhead through various compression mechanisms, such as with enhanced compressed RTP (ECRTP) [RFC3545] and robust header compression (ROHC) [RFC3095], and also to increase scalability of header compression. MPLS is used to route header-compressed (HC) packets over an MPLS label switched path (LSP) without compression/decompression cycles at each router. Such an HC over

MPLS capability can increase bandwidth efficiency as well as the processing scalability of the maximum number of simultaneous compressed flows that use header compression at each router.

Goals and requirements for header compression over MPLS are discussed in [RFC4247]. The solution put forth in this document using MPLS pseudowire technology has been designed to address these goals and requirements.

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 [RFC2119].

Forwarding Equivalence Class (FEC): a group of packets that are forwarded in the same manner (e.g., over the same LSP, with the same forwarding treatment)

Header Compression scheme (HC scheme): a particular method of performing header compression and its associated protocol. Multiple methods of header compression have been defined, including Robust Header Compression (ROHC [RFC3095]), compressed RTP (cRTP, [RFC2508]), enhanced cRTP (ECRTP, [RFC3545]), and IP Header Compression (IPHC, [RFC2507]). This draft explicitly supports all of the HC schemes listed above, and is intended to be extensible to others that may be developed.

Header Compression session (HC session): A session established between a header compressor and a header decompressor using a single HC scheme, over which multiple individual flows may be compressed. An HC session should not be confused with the individual traffic flows that may be compressed using a single Context ID. Each HC session manages a set of unique Context ID's.

Label: a short fixed length physically contiguous identifier which is used to identify a FEC, usually of local significance

Label Switched Path (LSP): the path through one or more LSRs at one level of the hierarchy followed by a packets in a particular forwarding equivalence class (FEC)

Label Switching Router (LSR): an MPLS node which is capable of forwarding native L3 packets label stack an ordered set of labels

MPLS domain: a contiguous set of nodes which operate MPLS routing and forwarding and which are also in one Routing or Administrative Domain

MPLS label: a label which is carried in a packet header, and which represents the packet's FEC

MPLS node: a node that is running MPLS. An MPLS node will be aware of MPLS control protocols, will operate one or more L3 routing

protocols, and will be capable of forwarding packets based on labels. An MPLS node may optionally be also capable of forwarding native L3 packets.

Multi Protocol Label Switching (MPLS): an IETF working group and the effort associated with the working group

Packet Switched Network (PSN): Within the context of PWE3, this is a network using IP or MPLS as the mechanism for packet forwarding.

Protocol Data Unit (PDU): The unit of data output to, or received from, the network by a protocol layer.

Pseudo Wire (PW): A mechanism that carries the essential elements of an emulated service from one provider edge router to one or more other provider edge routers over a PSN

Pseudo Wire Emulation Edge to Edge (PWE3): A mechanism that emulates the essential attributes of service (such as a T1 leased line or Frame Relay) over a PSN

Pseudo Wire PDU (PW-PDU): A PDU sent on the PW that contains all of the data and control information necessary to emulate the desired service

PSN Tunnel: A tunnel across a PSN, inside which one or more PWs can be carried

PSN Tunnel Signaling: Used to set up, maintain, and tear down the underlying PSN tunnel

PW Demultiplexer: Data-plane method of identifying a PW terminating at a provider edge router

Tunnel: A method of transparently carrying information over a network

3. Header Compression over MPLS Protocol Overview

To implement header compression (HC) over MPLS, after the ingress router applies the HC algorithm to the IP packet, the compressed packet is forwarded on an MPLS LSP using MPLS labels, and then the egress router restores the uncompressed header. Figure 1 illustrates an HC over MPLS session established on an LSP that traverses several label switch routers, from R1/HC --> R2 --> R3 --> R4/HD, where R1/HC is the ingress router performing header compression (HC), and R4/HD is the egress router performing header decompression (HD). This example assumes that the packet flow being compressed has RTP/UDP/IP headers and is using a HC scheme such as ROHC, cRTP or ECRTP. Compression of the RTP/UDP/IP header is performed at R1/HC, and the

compressed packets are routed using MPLS labels from R1/HC to R2, to R3, and finally to R4/HD, without further decompression/recompression

cycles. The RTP/UDP/IP header is decompressed at R4/HD and can be forwarded to other routers, as needed.

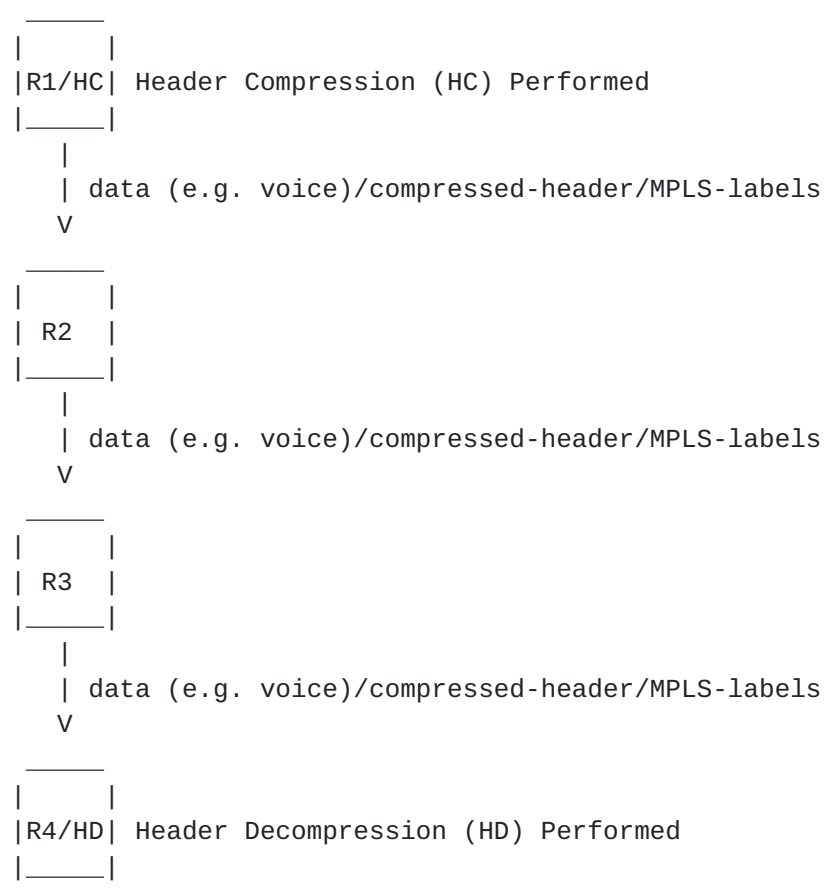

Figure 1. Example of HC over MPLS over Routers R1 --> R4

In the example scenario, header compression therefore takes place between R1 and R4, and the MPLS path transports data/compressed-header/MPLS-labels instead of data/RTP/UDP/IP/MPLS-labels, saving 36 octets per packet in the /RTP/UDP/IP/ header. Typically there are two MPLS labels (8 octets) and a link-layer (pseudowire) control word (2 octets). The MPLS label stack and link-layer headers are not compressed. Therefore HC over MPLS can significantly reduce the header overhead through compression mechanisms.

HC reduces the IP/UDP/RTP headers to 2-4 bytes for most packets. Half of the reduction in header size comes from the observation that half of the bytes in the IP/UDP/RTP headers remain constant over the life of the connection. After sending the uncompressed header template once, these fields may be removed from the compressed headers that follow. The remaining compression comes from the observation that although several fields change in every packet, the

difference from packet to packet is often constant and therefore the second-order difference is zero.

By maintaining both the uncompressed header and the first-order differences in the session state shared between the compressor and decompressor, all that must be communicated is an indication that the second-order difference was zero. In that case, the decompressor can reconstruct the original header without any loss of information simply by adding the first-order differences to the saved uncompressed header as each compressed packet is received. The compressed packet carries the context identification (CID), to indicate in which session context that packet should be interpreted.

MPLS pseudowires (PWs) [RFC3985] are used to transport the header compressed packets between the ingress and egress MPLS label switched router (LSR), and the PWs define one or more point-to-point instances corresponding to each HC session at the header decompressor. Compressed data is routed on a separate MPLS LSP/PW from compressor to decompressor. The decompressor uses the incoming MPLS PW Label and the CID to locate the proper decompression context. Standard HC methods (e.g., ECRTP, ROHC, etc.) are used to determine the context. The CIDs are assigned by the HC as normal, and there would be no problem if duplicate CIDs are received at the HD for different compressed sessions. For example, if HCa and HCb assign the same CID to each of 2 separate flows, each PW then had a logically separate HD instance, in this case, defined by the <PWlabel-HCa, CID> <PWlabel-HCb, CID> tuples, independent of all other PWs. That is, HCa and HCb have a separate decompression context for the two flows based on the PW label and CID mapping.

An MPLS PW allows protocol data units for various link-layer protocols to be encapsulated and carried over an MPLS network. In this approach, compressed packets are encapsulated and transported over a PW across the MPLS network using MPLS labels, which include the packet switched network (PSN) label and PW label. A PW control word is used to identify the type of packet, a unique PW Type is defined for each HC scheme, and, as normal, a CID is used to identify each compressed packet context and payload. Each HC scheme is applied directly over its own PW type. The PW is set up by the PW signaling protocol [PW-SIG], and messages for HC session setup and HC parameter negotiation [RFC3241, RFC3544] are reused to enable HC session configuration

Figure 1 illustrates an example data flow set up from R1/HC --> R2 --> R3 --> R4/HD, where R1/HC is the ingress router where header compression is performed, and R4/HD is the egress router where header decompression is done. Each router functions as an LSR and supports signaling of LSP/PWs. A summary of the procedures is as follows:

1. [PW-SIG] is used to create the R1 --> R4 LSP/PW that follows R1 --> R2 --> R3 --> R4.

2. [PW-SIG] is used to create the R4 --> R1 LSP/PW that follows R4 --> R3 --> R2 --> R1.
3. [RFC3544] and [RFC3246] are used to negotiate HC scheme

parameters, which is extended in this specification to negotiating during PW setup, as specified in Section 4.1.
4. R1/HC assigns a CID to the flow and uses the R1 --> R4 LSP/PW to send HC scheme control packets and compressed packets to R4/HC, with LSP and PW labels.
5. R4/HD uses the incoming MPLS PW label and CID to locate the proper decompression context to decompress the compressed packets sent by R1/HC.
6. R4/HC assigns a CID to the flow and uses the R4 --> R1 LSP/PW to send HC scheme control packets and compressed packets to R1/HD, with LSP and PW labels.
7. R1/HD uses the incoming MPLS PW label and CID to locate the proper decompression context to decompress the compressed packets sent by R4/HC.
8. if needed to resync, R4/HD sends an appropriate HC scheme control packet to R1/HC; R1/HC responds with the appropriate HC scheme control packet to R4/HD.
9. if needed to resync, R1/HD sends an appropriate HC scheme control packet to R4/HC; R4/HC responds with the HC scheme control packet to R1/HD.
10. Existing HC scheme procedures are used to assign and free up the CIDs; see, for example, Section 7 in [ROHC-IMPL-GUIDE].

4. Protocol Specifications

Figure 2 illustrates the PW stack reference model to support PW emulated services.

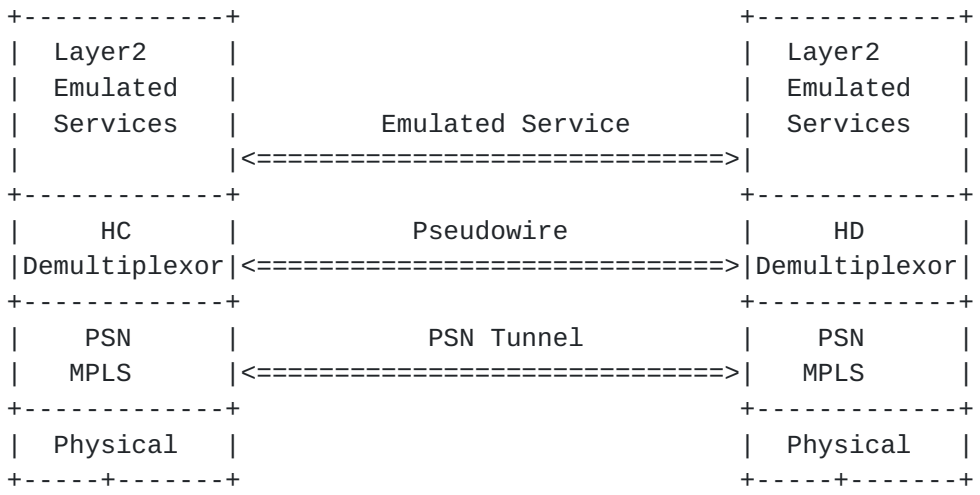

Figure 2: Pseudowire Protocol Stack Reference Model

Each HC-HD compressed session MUST be identified by the PW label. A single PW label MUST be used for many HC flows (could be 100's or 1000's) rather than assigning a different PW label to each flow. The

latter approach would involve a complex mechanism for PW label assignment, freeing up of labels after a flow terminates, etc., for potentially 1000's of simultaneous HC flows. On the other hand, the

mechanism for CID assignment, freeing up, etc. is in place and there is no need to duplicate it with PW assignment/deassignment for individual HC flows.

Multiple PWs SHOULD be established in case different QoS requirements are needed for different compressed streams. The QoS received by the flow would be determined by the EXP bit marking in the PW label. Normally, all RTP packets would get the same EXP marking, equivalent to EF treatment in DiffServ. However, the protocol specified in this document applies to several different types of streams, not just RTP streams, and QoS treatment other than EF may be required for those streams.

Figure 3 shows the HC over MPLS protocol stack (with uncompressed header):

Media stream
RTP
UDP
IP
PW control word
MPLS label stack (at least 2 labels for this application)
Link layer under MPLS (PPP, PoS, Ethernet)
Physical layer (SONET/SDH, fiber, copper)

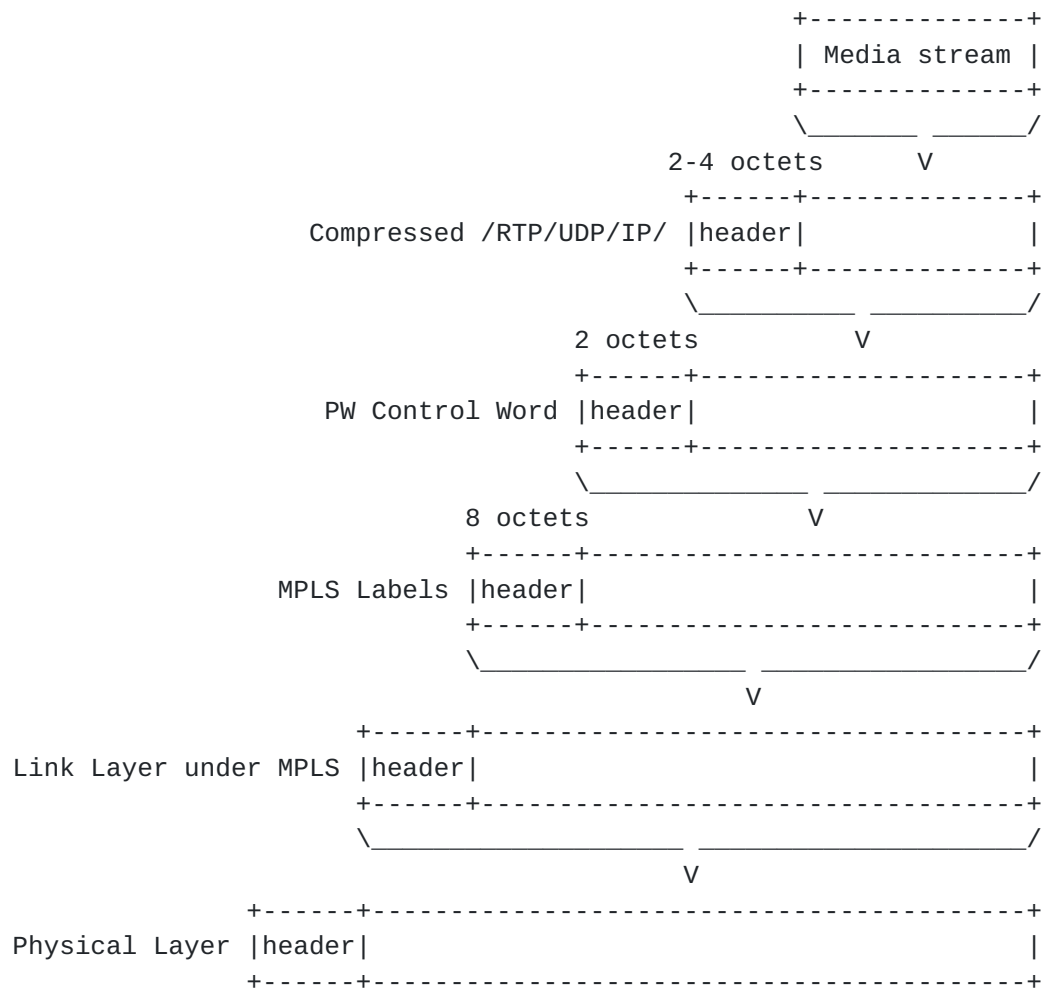

Figure 3 - Header Compression over MPLS Media Stream Transport

The PW control word MUST be to used to identify the packet types for the HC scheme in use. The MPLS labels technically define two layers: the PW identifier and the MPLS tunnel identifier. The PW label MUST be used as the demultiplexer field by the HD, where the PW label appears at the bottom label of an MPLS label stack. There can also be other MPLS labels, for example, to identify an MPLS VPN. The IP/UDP/RTP headers are compressed before transmission, leaving the rest of the stack alone, as shown in Figure 3.

4.1 MPLS Pseudowire Setup & Signaling

PWs MUST be set up in advance for the transport of media streams using [PW-SIG] control messages exchanged by the HC-HD endpoints. Furthermore, a PW type MUST be used to indicate the HC scheme being used on the PW. [PW-SIG] specifies the MPLS label distribution protocol (LDP) [RFC3036] extensions to set up and maintain the PWs, and defines new LDP objects to identify and signal attributes of PWs.

Any acceptable method of MPLS label distribution MAY be used for distributing the MPLS tunnel label. To assign and distribute the PW labels, an LDP session MUST be set up between the PW endpoints using

the extended discovery mechanism described in [RFC3036]. The PW label bindings are distributed using the LDP downstream unsolicited mode described in [RFC3036]. An LDP label mapping message contains a forward equivalence class (FEC) object, a label object, and possible other optional objects. The FEC object indicates the meaning of the label, identifies the PW type, and identifies the PW that the PW label is bound to. See [PW-SIG] for further explanation of PW signaling.

This specification defines new PW type values to be carried within the FEC object to identify HC PWs for each HC scheme. The PW type is a 15-bit parameter assigned by IANA, as specified in the [IANA] registry, and MUST be used to indicate the HC scheme being used on the PW. The following PW type values have been reserved [IANA]:

0x001A	ROHC Transport Header-compressed Packets	[RFC3095]
0x001B	eCRTP Transport Header-compressed Packets	[RFC3545]
0x001C	IPHC Transport Header-compressed Packets	[RFC2507]
0x001D	cRTP Transport Header-compressed Packets	[RFC2508]

The PW control word enables distinguishing between various packets types (e.g., uncompressed, UDP compressed, RTP compressed, context-state, etc.). However, the PW control word indications are not unique across HC schemes, and therefore the PW type value allows the HC scheme to be identified.

4.2 Header Compression Scheme Setup, Negotiation, & Signaling

As described in the previous section, the HC PW MUST be used for compressed packets only, which is configured at PW setup. If a flow is not compressed, it MUST NOT ne placed on the HC PW. HC scheme parameters MAY be configured, or the Interface Parameters Sub-TLV MUST be used if the setup parameters are signaled, as now described.

The PW HC approach relies on the PW/MPLS layer to convey HC session configuration information. The Interface Parameters Sub-TLV [IANA, PW-SIG] MUST be used to signal HC session setup and HC parameter negotiation using the TLVs/messages described in [RFC3241, RFC3544]. That is, the messages specified in [RFC3241, RFC3544] are reused in this specification to specify PW specific parameters, and to configure the HC and HD ports at the edges of the PW, so that they have the necessary capabilities to interoperate with each other.

Pseudowire Interface Parameter Sub-TLV type values are specified in [IANA]. Two code-points have been reserved, as follows:

Parameter ID	Length	Description	References
0x0D	up to 256 bytes	ROHC over MPLS configuration	RFC 3241
0x0F	up to 256 bytes	CRTP/ECRTP/IPHC HC over MPLS	RFC 3544

configuration

TLVs identified in [RFC3241] and [RFC3544] MUST be encapsulated in the PW Interface Parameters Sub-TLV and used to negotiate header compression session setup and parameter negotiation for their respective protocols. The TLVs supported in this manner MUST include the following:

- o Configuration Option Format, RTP-Compression Suboption, Enhanced RTP-Compression Suboption, TCP/non-TCP Compression Suboptions, as specified in [RFC3544]
- o Configuration Option Format, PROFILES Suboption, as specified in [RFC3241]

These TLVs are now specified in the following sections.

4.2.1 Configuration Option Format [RFC3544]

Both the network control protocol for IPv4, IPCP [RFC1332] and the IPv6 NCP, IPV6CP [RFC2472] may be used to negotiate IP Header Compression parameters for their respective protocols. The format of the configuration option is the same for both IPCP and IPV6CP.

Description

This NCP configuration option is used to negotiate parameters for IP Header Compression. Successful negotiation of parameters enables the use of Protocol Identifiers FULL_HEADER, COMPRESSED_TCP, COMPRESSED_TCP_NODELTA, COMPRESSED_NON_TCP and CONTEXT_STATE as specified in [RFC2507]. The option format is summarized below. The fields are transmitted from left to right.

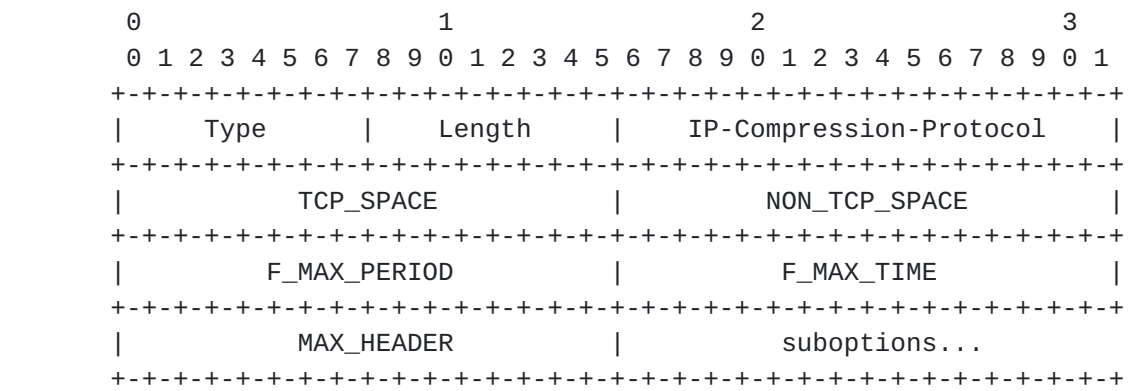

Type
2

Length
>= 14

The length may be increased if the presence of additional

parameters is indicated by additional suboptions.

IP-Compression-Protocol
0061 (hex)

TCP_SPACE
The TCP_SPACE field is two octets and indicates the maximum value of a context identifier in the space of context identifiers allocated for TCP.

Suggested value: 15

TCP_SPACE must be at least 0 and at most 255 (the value 0 implies having one context).

NON_TCP_SPACE

The NON_TCP_SPACE field is two octets and indicates the maximum value of a context identifier in the space of context identifiers allocated for non-TCP. These context identifiers are carried in COMPRESSED_NON_TCP, COMPRESSED_UDP and COMPRESSED_RTP packet headers.

Suggested value: 15

NON_TCP_SPACE must be at least 0 and at most 65535 (the value 0 implies having one context).

F_MAX_PERIOD

Maximum interval between full headers. No more than F_MAX_PERIOD COMPRESSED_NON_TCP headers may be sent between FULL_HEADER headers.

Suggested value: 256

A value of zero implies infinity, i.e. there is no limit to the number of consecutive COMPRESSED_NON_TCP headers.

F_MAX_TIME

Maximum time interval between full headers. COMPRESSED_NON_TCP headers may not be sent more than F_MAX_TIME seconds after sending the last FULL_HEADER header.

Suggested value: 5 seconds

A value of zero implies infinity.

MAX_HEADER

The largest header size in octets that may be compressed.

Suggested value: 168 octets

The value of MAX_HEADER should be large enough so that at least the outer network layer header can be compressed. To increase compression efficiency MAX_HEADER should be set to a value large enough to cover common combinations of network and transport layer headers.

suboptions

The suboptions field consists of zero or more suboptions. Each suboption consists of a type field, a length field and zero or more parameter octets, as defined by the suboption type. The

value of the length field indicates the length of the suboption in its entirety, including the lengths of the type and length fields.

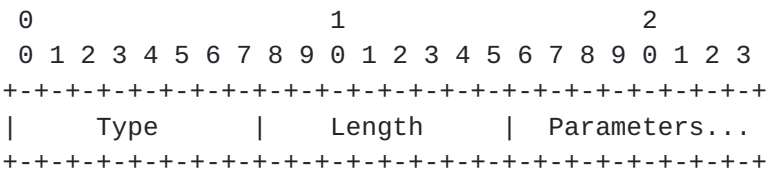

4.2.2 RTP-Compression Suboption [RFC3544]

The RTP-Compression suboption is included in the NCP IP-Compression-Protocol option for IPHC if IP/UDP/RTP compression is to be enabled.

Inclusion of the RTP-Compression suboption enables use of additional Protocol Identifiers COMPRESSED_RTP and COMPRESSED_UDP along with additional forms of CONTEXT_STATE as specified in [RFC2508].

Description

Enable use of Protocol Identifiers COMPRESSED_RTP, COMPRESSED_UDP and CONTEXT_STATE as specified in [RFC2508].

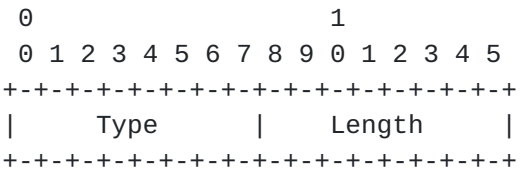

Type
1

Length
2

4.2.3 Enhanced RTP-Compression Suboption [RFC3544]

To use the enhanced RTP header compression defined in [RFC3545], a new sub-option 2 is added. Sub-option 2 is negotiated instead of, not in addition to, sub-option 1.

Description

Enable use of Protocol Identifiers COMPRESSED_RTP and CONTEXT_STATE as specified in [RFC2508]. In addition, enable use of [RFC3545] compliant compression including the use of Protocol Identifier COMPRESSED_UDP with additional flags and use of the C flag with the FULL_HEADER Protocol Identifier to indicate use of HDRCKSUM with COMPRESSED_RTP and COMPRESSED_UDP packets.

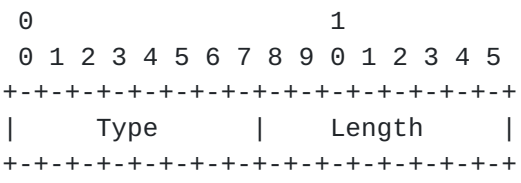

Type
2

Length
2

4.2.4 Negotiating header compression for only TCP or only non-TCP Packets [RFC3544]

In RFC 2509 it was not possible to negotiate only TCP header compression or only non-TCP header compression because a value of 0 in the TCP_SPACE or the NON_TCP_SPACE fields actually means that 1 context is negotiated.

A new suboption 3 is added to allow specifying that the number of contexts for TCP_SPACE or NON_TCP_SPACE is zero, disabling use of the corresponding compression.

Description

Enable header compression for only TCP or only non-TCP packets.

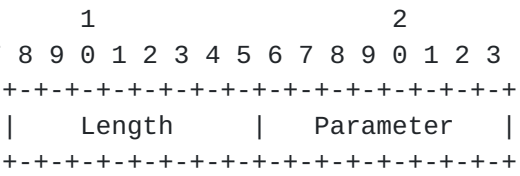

Type
3

Length
3

Parameter

The parameter is 1 byte with one of the following values:

1 = the number of contexts for TCP_SPACE is 0
2 = the number of contexts for NON_TCP_SPACE is 0

This suboption overrides the values that were previously assigned to TCP_SPACE and NON_TCP_SPACE in the IP Header Compression option.

If suboption 3 is included multiple times with parameter 1 and 2,

compression is disabled for all packets.

4.2.5 Configuration Option Format [RFC3241]

Both the network control protocol for IPv4, IPCP [RFC1332] and the IPv6 NCP, IPV6CP [RFC2472] may be used to negotiate IP Header Compression parameters for their respective protocols. The format of the configuration option is the same for both IPCP and IPV6CP.

Description

This NCP configuration option is used to negotiate parameters for Robust Header Compression. The option format is summarized below. The fields are transmitted from left to right.

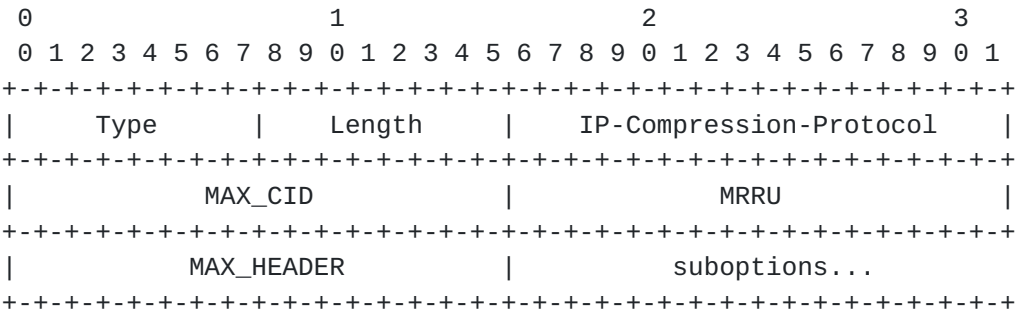

Type
2

Length
>= 10

The length may be increased if the presence of additional parameters is indicated by additional suboptions.

IP-Compression-Protocol
0003 (hex)

MAX_CID
The MAX_CID field is two octets and indicates the maximum value of a context identifier.

Suggested value: 15

MAX_CID must be at least 0 and at most 16383 (The value 0 implies having one context).

MRRU
The MRRU field is two octets and indicates the maximum reconstructed reception unit (see [RFC3095], section 5.1.1).

Suggested value: 0

MAX_HEADER
The largest header size in octets that may be compressed.

Suggested value: 168 octets

The value of MAX_HEADER should be large enough so that at least the outer network layer header can be compressed. To increase compression efficiency MAX_HEADER should be set to a value large enough to cover common combinations of network and transport layer headers.

NOTE: The four ROHC profiles defined in RFC 3095 do not provide for a MAX_HEADER parameter. The parameter MAX_HEADER defined by this document is therefore without consequence in these profiles. Other profiles (e.g., ones based on RFC 2507) can make use of the parameter by explicitly referencing it.

suboptions
The suboptions field consists of zero or more suboptions. Each suboption consists of a type field, a length field and zero or more parameter octets, as defined by the suboption type. The value of the length field indicates the length of the suboption in its entirety, including the lengths of the type and length fields.

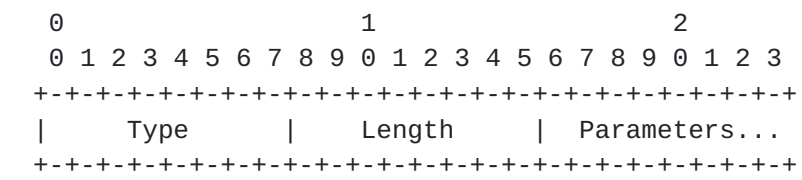

4.2.6 PROFILES Suboption [RFC3241]

The set of profiles to be enabled is subject to negotiation. Most initial implementations of ROHC implement profiles 0x0000 to 0x0003. This option MUST be supplied.

Description

Define the set of profiles supported by the decompressor.

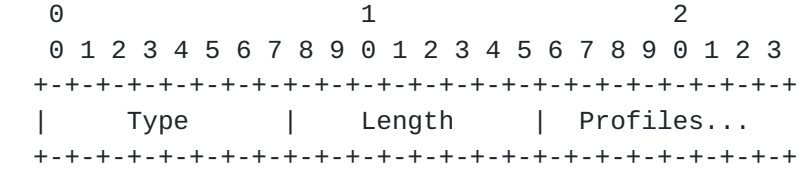

Type
1

Length
2n+2

Value
n octet-pairs in ascending order, each octet-pair specifying a ROHC profile supported.

HC flow identification is being done now in many ways. Since there are multiple possible approaches to the problem, no specific method is specified in this document.

4.3 Encapsulation of Header Compressed Packets

The PW control word is used to identify the packet types for IPHC [RFC2507], cRTP [RFC2508], and ECRTP [RFC3545], as shown in Figure 4:

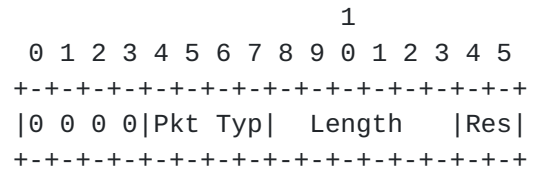

Figure 4 - PW Control Word

where:

"Packet Type" encoding:
0: ROHC Small-CIDs
1: ROHC Large-CIDs
2: FULL_HEADER
3: COMPRESSED_TCP
4: COMPRESSED_TCP_NODELTA
5: COMPRESSED_NON_TCP
6: COMPRESSED_RTP_8
7: COMPRESSED_RTP_16
8: COMPRESSED_UDP_8
9: COMPRESSED_UDP_16
10: CONTEXT_STATE
11-15: TO BE ASSIGNED BY IANA (see Section 7, IANA considerations, for discussion of the Registry)

As discussed in [ECMP-AVOID], since this MPLS payload type is not IP, the first nibble is set to 0000 to avoid being mistaken for IP. This is also consistent with the encoding of the PWE3 control word [PW-CNTL-WORD].

Note that ROHC [RFC3095] provides its own packet type within the protocol, however the PW control word MUST still be used to avoid the problems identified above. For ROHC, the first byte of the PW

control word is set to zero.

The PW control word length field is ONLY used for short packets because padding may be appended by the Ethernet Data Link Layer. If the length is >= than 64 octets, the length field MUST be set to zero [PW-CNTL-WORD]. If the MPLS payload is less than 64 bytes, the length field MUST be set to the length of the PW payload plus the length of the PW control word. Note that the last 2 bits in the PW control word are reserved.

4.4 Packet Reordering

Packet reordering for ROHC and ECRTP are discussed in [ROHC-REORDER] and [ECRTP-REORDER], which are a useful source of information. In case of lossy links and other reasons for reordering, implementation adaptations are needed to allow all the schemes to be used in this case. CRTP is viewed as having risks for a number PW environments due to misordering and loss, although commercial issues lead to its choice. In these circumstances, it must be implemented and deployed with care. IPHC should use TCP_NODELTA, ECRTP should send absolute values, ROHC should be adapted as discussed in [ROHC-REORDER], and ECRTP should be adapted as discussed in [ECRTP-REORDER]. An evaluation and simulation of ECRTP and ROHC reordering is given in [REORDER-EVAL].

5. Security Considerations

MPLS pseudowire security considerations in general are discussed in [RFC3985] and [PW-SIG], and those considerations also apply to this document. This document specifies an encapsulation and not the protocols that may be used to carry the encapsulated packets across the PSN, or the protocols being encapsulated. Each such protocol may have its own set of security issues, but those issues are not affected by the encapsulations specified herein.

6. Acknowledgements

The authors appreciate valuable inputs and suggestions from Loa Andersson, Scott Brim, Stewart Bryant, Adrian Farrel, Victoria Fineberg, Allison Mankin, Luca Martini, Colin Perkins, Kristofer Sandlund, Yaakov Stein, George Swallow, Curtis Villamizar, and Magnus Westerlund.

7. IANA Considerations

As discussed in Section 4.1, PW type values need to be assigned by IANA, as follows:

0x001A	ROHC Transport Header-compressed Packets	[RFC3095]
0x001B	eCRTP Transport Header-compressed Packets	[RFC3545]
0x001C	IPHC Transport Header-compressed Packets	[RFC2507]

0x001D cRTP Transport Header-compressed Packets   [RFC2508]

As discussed in Section 4.2, Pseudowire Interface Parameter Sub-TLV type values need to be specified by IANA, as follows:

Parameter ID	Length	Description	References
0x0D	up to 256 bytes	ROHC over MPLS configuration	RFC 3241
0x0F	up to 256 bytes	CRTP/ECRTP/IPHC HC over MPLS configuration	RFC 3544

As discussed in Section 4.3, IANA needs to define a new registry, "Header Compression Over MPLS PW Control Word Packet Type". This is a four bit value. Packet Types 0 through 10 are defined in Section 4.3 of this document. Packet Types 11 to 15 are to be assigned by IANA using the "Expert Review" policy defined in [RFC2434].

8. Normative References

[PW-SIG] Martini, L., et. al., "Pseudowire Setup and Maintenance Using LDP," work in progress.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119, March 1997.

[RFC2434] Narten, T. et al, "Guidelines for Writing an IANA Considerations Section in RFCs", RFC 2434, BCP 26, October 1998.

[RFC3036] Andersson, L., et. al., "LDP Specification," RFC 3036, January 2001.

[RFC3241] Bormann, C., "Robust Header Compression (ROHC) over PPP," RFC 3241, April 2002.

[RFC3544] Engan, M., Casner, S., Bormann, C., "IP Header Compression over PPP", RFC 3544, July 2003.

[RFC3985] Bryant, S., Pate, P., "Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture," RFC 3985, March 2005.

9. Informative References

[ECMP-AVOID] Swallow, G., et. al., "Avoiding Equal Cost Multipath Treatment in MPLS Networks," work in progress.

[ECRTP-REORDER] Koren, T., et. al., "Packet reordering in Extended CRTP (ECRTP)," work in progress.

[IANA] Martini, L., et. al., "IANA Allocations for Pseudo Wire Edge To Edge Emulation (PWE3)," work in progress.

[PW-CNTL-WORD] Bryant, S., et. al., "PWE3 Control Word for use over

an MPLS PSN," work in progress.

[PW-PPP] Martini, L., et. al., "Encapsulation Methods for Transport of PPP/HDLC Over IP and MPLS Networks," work in progress.

[REORDER-EVAL] Knutsson, C., "Evaluation and Implementation of Header Compression Algorithm ECRTP," http://epubl.luth.se/1402-1617/2004/286/LTU-EX-04286-SE.pdf.

[RFC1332] McGregor, G., "The PPP Internet Protocol Control Protocol (IPCP)," May 1992.

[RFC2507] Degermark, M., et. al., "IP Header Compression," RFC 2507, February 1999.

[RFC2508] Casner, S., Jacobsen, V., "Compressing IP/UDP/RTP Headers for Low-Speed Serial Links", RFC 2508, February 1999.

[RFC2547] Rosen, E., Rekhter, Y., "BGP/MPLS VPNs", RFC 2547, March 1999.

[RFC3095] Bormann, C., et. al., "RObust Header Compression (ROHC): Framework and four profiles: RTP, UDP, ESP, and uncompressed," RFC 3095, July 2001.

[RFC3545] Koren, T., et. al., "Compressing IP/UDP/RTP Headers on Links with High Delay, Packet Loss, and Reordering," RFC 3545, July 2003.

[RFC4247] Ash, G., Goode, B., Hand, J., "Requirements for Header Compression over MPLS", RFC 4247, November 2005.

[ROHC-IMPL-GUIDE] Jonsson, L-E., Kremer, P. "The RFC 3095 Implementer's Guide," work in progress.

[ROHC-REORDER] Pellitier, G., et. al., "RObust Header Compression (ROHC): ROHC over Channels that can Reorder Packets," work in progress.

10. Authors' & Contributors' Addresses

Jerry Ash (Editor)
AT&T
Room MT D5-2A01
200 Laurel Avenue
Middletown, NJ 07748, USA
Phone: +1 732-420-4578
Email: gash@att.com

Andrew G. Malis (Editor)
Tellabs

90 Rio Robles Dr.
San Jose, CA 95134

Email: Andy.Malis@tellabs.com

Bur Goode
AT&T
Phone: +1 203-341-8705
Email: bgoode@att.com

Jim Hand
AT&T
Room MT A2-4F36
200 Laurel Avenue
Middletown, NJ 07748, USA
Phone: +1 732-420-6179
Email: jameshand@att.com

Lars-Erik Jonsson
Ericsson AB
Box 920
SE-971 28 Lulea, Sweden
Phone: +46 8 404 29 61
EMail: lars-erik.jonsson@ericsson.com

Raymond Zhang
Infonet Services Corporation
2160 E. Grand Ave. El Segundo, CA 90025 USA
Email: zhangr@bt.infonet.com

Intellectual Property Statement

ietf-ipr@ietf.org.

Disclaimer of Validity

Copyright Statement