             Use of Label Switching on Frame Relay Networks
                              Specification

Status of this Memo

Copyright Notice

Abstract

   This document defines the model and generic mechanisms for
   Multiprotocol Label Switching on Frame Relay networks. Furthermore,
   it extends and clarifies portions of the Multiprotocol Label
   Switching Architecture described in [ARCH] and the Label Distribution
   Protocol (LDP) described in [LDP] relative to Frame Relay Networks.
   MPLS enables the use of Frame Relay Switches as Label Switching
   Routers (LSRs).

Table of Contents

1. Introduction

   The Multiprotocol Label Switching Architecture is described in
   [ARCH]. It is possible to use Frame Relay switches as Label
   Switching Routers. Such Frame Relay switches run network layer
   routing algorithms (such as OSPF, IS-IS, etc.), and their forwarding
   is based on the results of these routing algorithms. No specific
   Frame Relay routing is needed.

   When a Frame Relay switch is used for label switching, the top
   (current) label, on which forwarding decisions are based, is carried
   in the DLCI field of the Frame Relay data link layer header of a
   frame. Additional information carried along with the top (current)
   label, but not processed by Frame Relay switching, along with other
   labels, if the packet is multiply labeled, are carried in the generic
   MPLS encapsulation defined in [STACK].

   Frame Relay permanent virtual circuits (PVCs) could be configured to
   carry label switching based traffic. The DLCIs would be used as MPLS
   Labels and the Frame Relay switches would become Frame Relay Label
   Switching Routers, while the MPLS traffic would be encapsulated
   according to this specification, and would be forwarded based on
   network layer routing information.

   The keywords MUST, MUST NOT, MAY, OPTIONAL, REQUIRED, RECOMMENDED,
   SHALL, SHALL NOT, SHOULD, SHOULD NOT are to be interpreted as defined
   in RFC 2119.

   This document is a companion document to [STACK] and [ATM].

2. Terminology

   LSR

      A Label Switching Router (LSR) is a device which implements the
      label switching control and forwarding components described in
      [ARCH].

   LC-FR

      A label switching controlled Frame Relay (LC-FR) interface is a
      Frame Relay interface controlled by the label switching control
      component. Packets traversing such an interface carry labels in
      the DLCI field.

   FR-LSR

      A FR-LSR is an LSR with one or more LC-FR interfaces which
      forwards frames between two such interfaces using labels carried
      in the DLCI field.

   FR-LSR domain

      A FR-LSR domain is a set of FR-LSRs, which are mutually
      interconnected by LC-FR interfaces.

   Edge Set

      The Edge Set of an FR-LSR domain is the set of LSRs, which are
      connected to the domain by LC-FR interfaces.

   Forwarding Encapsulation

      The Forwarding Encapsulation is the type of MPLS encapsulation
      (Frame Relay, ATM, Generic) of a packet that determines the
      packet's MPLS forwarding, or the network layer encapsulation if
      that packet is forwarded based on the network layer (IP,
      etc...)header.

   Input Encapsulation

      The Input Encapsulation is the type of MPLS encapsulation (Frame
      Relay, ATM, Generic) of a packet when that packet is received on
      an LSR's interface, or the network layer (IP, etc...)encapsulation
      if that packet has no MPLS encapsulation.

   Output Encapsulation

      The Output Encapsulation is the type of MPLS encapsulation (Frame
      Relay, ATM, Generic) of a packet when that packet is transmitted
      on an LSR's interface, or the network layer (IP,
      etc...)encapsulation if that packet has no MPLS encapsulation.

   Input TTL

      The Input TTL is the MPLS TTL of the top of the stack when a
      labeled packet is received on an LSR interface, or the network
      layer (IP) TTL if the packet is not labeled.

   Output TTL

      The Output TTL is the MPLS TTL of the top of the stack when a
      labeled packet is transmitted on an LSR interface, or the network
      layer (IP) TTL if the packet is not labeled.

   Additionally, this document uses terminology from [ARCH].

3. Special characteristics of Frame Relay Switches

   While the label switching architecture permits considerable
   flexibility in LSR implementation, a FR-LSR is constrained by the
   capabilities of the (possibly pre-existing) hardware and the
   restrictions on such matters as frame format imposed by the
   Multiprotocol Interconnect over Frame Relay [MIFR], or Frame Relay
   standards [FRF], etc.... Because of these constraints, some special
   procedures are required for FR-LSRs.

   Some of the key features of Frame Relay switches that affect their
   behavior as LSRs are:

   - the label swapping function is performed on fields (DLCI) in the
      frame's Frame Relay data link header; this dictates the size and
      placement of the label(s) in a packet. The size of the DLCI field
      can be 10 (default) or 23 bits, and it can span two or four bytes
      in the header.

   - there is generally no capability to perform a 'TTL-decrement'
      function as is performed on IP headers in routers.

   - congestion control is performed by each node based on parameters
      that are passed at circuit creation. Flags in the frame headers
      may be set as a consequence of congestion, or exceeding the
      contractual parameters of the circuit.

     - although in a standard switch it may be possible to configure
        multiple input DLCIs to one output DLCI resulting in a
        multipoint-to-point circuit, multipoint-to-multipoint VCs are
        generally not fully supported.

     This document describes ways of applying label switching to Frame
     Relay switches, which work within these constraints.

4. Label Encapsulation

     By default, all labeled packets should be transmitted with the
     generic label encapsulation as defined in [STACK], using the frame
     relay null encapsulation mechanism:

```
                 0                       1               (Octets)
                +-----------------------+----------------------+
   (Octets)0    |                       |                      |
                /         Q.922 Address                        /
                /       (length 'n' equals 2 or 4)             /
                |                       |                      |
                +---------------------+----------------------+
         n      |                     .                      |
                /                     .                      /
                /         MPLS packet                        /
                |                     .                      |
                +-----------------------+----------------------+
```
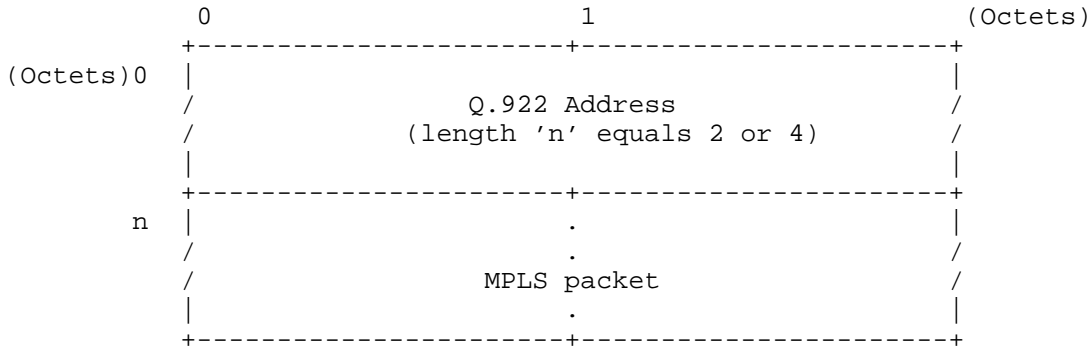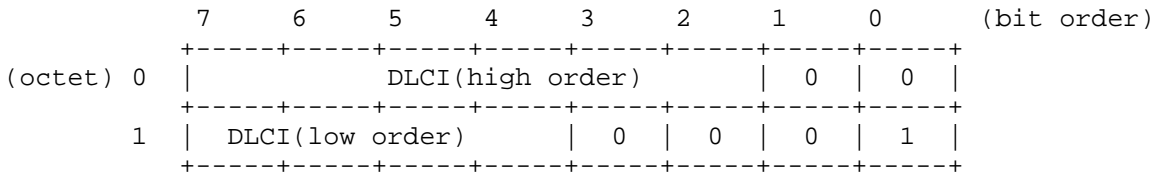
     "n" is the length of the Q.922 Address which can be 2 or 4 octets.

     The Q.922 [ITU] representation of a DLCI (in canonical order -
     the first bit is stored in the least significant, i.e., the
     right-most bit of a byte in memory) [CANON] is the following:

```
             7     6     5     4     3     2     1     0    (bit order)
            +-----+-----+-----+-----+-----+-----+-----+-----+
(octet) 0  |           DLCI(high order)      | 0   | 0   |
            +-----+-----+-----+-----+-----+-----+-----+-----+
        1  |   DLCI(low order)       | 0   | 0   | 0   | 1   |
            +-----+-----+-----+-----+-----+-----+-----+-----+

               10 bits DLCI
```
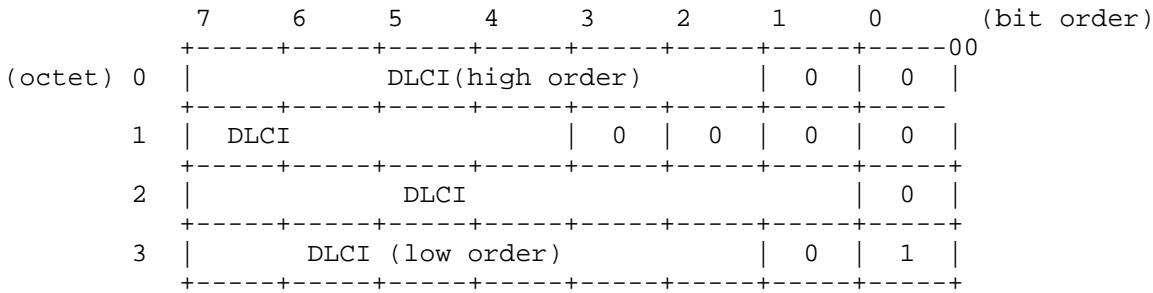
```
            7     6     5     4     3     2     1     0    (bit order)
          +-----+-----+-----+-----+-----+-----+-----+-----00
(octet) 0 |           DLCI(high order)       | 0   | 0   |
          +-----+-----+-----+-----+-----+-----+-----+-----
        1 | DLCI                 | 0   | 0   | 0   | 0   |
          +-----+-----+-----+-----+-----+-----+-----+-----+
        2 |             DLCI                       | 0   |
          +-----+-----+-----+-----+-----+-----+-----+-----+
        3 |         DLCI (low order)         | 0   | 1   |
          +-----+-----+-----+-----+-----+-----+-----+-----+

            23 bits DLCI
```

The use of the frame relay null encapsulation implies that labels
implicitly encode the network protocol type.

Rules regarding the construction of the label stack, and error
messages returned to the frame source are also described in [STACK].

The generic encapsulation contains "n" labels for a label stack of
depth "n" [STACK], where the top stack entry carries significant
values for the EXP, S , and TTL fields [STACK] but not for the label,
which is rather carried in the DLCI field of the Frame Relay data
link header encoded in Q.922 [ITU] address format.

5. Frame Relay Label Switching Processing

5.1 Use of DLCIs

Label switching is accomplished by associating labels with routes and
using the label value to forward packets, including determining the
value of any replacement label. See [ARCH] for further details. In
a FR-LSR, the top (current) MPLS label is carried in the DLCI field
of the Frame Relay data link layer header of the frame. The top
label carries implicitly information about the network protocol type.

For two connected FR-LSRs, a full-duplex connection must be available
for LDP. The DLCI for the LDP VC is assigned a value by way of
configuration, similar to configuring the DLCI used to run IP routing
protocols between the switches.

With the exception of this configured value, the DLCI values used for
MPLS in the two directions of the link may be treated as belonging to
two independent spaces, i.e., VCs may be half-duplex, each direction
with its own DLCI.

The allowable ranges of DLCIs, the size of DLCIs, and the support for
VC merging MUST be communicated through LDP messages. Note that the
range of DLCIs used for labels depends on the size of the DLCI field.

5.2 Homogeneous LSPs

If <LSR1, LSR2, LSR3> is an LSP, it is possible that LSR1, LSR2, and
LSR3 will use the same encoding of the label stack when transmitting
packet P from LSR1, to LSR2, and then to LSR3. Such an LSP is
homogeneous.

5.3 Heterogeneous LSPs

If <LSR1, LSR2, LSR3> is an LSP, it is possible that LSR1 will use
one encoding of the label stack when transmitting packet P to LSR2,
but LSR2 will use a different encoding when transmitting a packet P
to LSR3. In general, the MPLS architecture supports LSPs with
different label stack encodings on different hops. When a labeled
packet is received, the LSR must decode it to determine the current
value of the label stack, then must operate on the label stack to
determine the new label value of the stack, and then encode the new
value appropriately before transmitting the labeled packet to its
next hop.

Naturally there will be MPLS networks which contain a combination of
Frame Relay switches operating as LSRs, and other LSRs, which operate
using other MPLS encapsulations, such as the Generic (MPLS shim
header), or ATM encapsulation. In such networks there may be some
LSRs, which have Frame Relay interfaces as well as MPLS Generic
("MPLS Shim") interfaces. This is one example of an LSR with
different label stack encodings on different hops of the same LSP.
Such an LSR may swap off a Frame Relay encoded label on an incoming
interface and replace it with a label encoded into a Generic MPLS
(MPLS shim) header on the outgoing interface.

5.4 Frame Relay Label Switching Loop Prevention and Control

FR-LSRs SHOULD operate on loop free FR-LSPs or LSP Frame Relay
segments. Therefore, FR-LSRs SHOULD use loop detection and MAY use
loop prevention mechanisms as described in [ARCH], and [LDP].

5.4.1 FR-LSRs Loop Control - MPLS TTL processing

The MPLS TTL encoded in the MPLS label stack is a mechanism used to:

(a) suppress loops;

(b) limit the scope of a packet.

When a packet travels along an LSP, it should emerge with the same
TTL value that it would have had if it had traversed the same
sequence of routers without having been label switched. If the
packet travels along a hierarchy of LSPs, the total number of LSR-
hops traversed should be reflected in its TTL value when it emerges
from the hierarchy of LSPs [ARCH].

The initial value of the MPLS TTL is loaded into a newly pushed label
stack entry from the previous TTL value, whether that is from the
network layer header when no previous label stack existed, or from a
pre-existent lower level label stack entry.

A FR-LSR switching same level labeled packets does not decrement the
MPLS TTL. A sequence of such FR-LSR is a "non-TTL segment".

When a packet emerges from a "non-TTL LSP segment", it should however
reflect in the TTL the number of LSR-hops it traversed. In the
unicast case, this can be achieved by propagating a meaningful LSP
length or LSP Frame Relay segment length to the FR-LSR ingress nodes,
enabling the ingress to decrement the TTL value before forwarding
packets into a non-TTL LSP segment [ARCH].

When an ingress FR-LSR determines upon decrementing the MPLS TTL that
a particular packet's TTL will expire before the packet reaches the
egress of the "non-TTL LSP segment", the FR-LSR MUST not label switch
the packet, but rather follow the specifications in [STACK] in an
attempt to return an error message to the packet's source:

    - it treats the packet as an expired packet and return an ICMP
       message to its source.

    - it forwards the packet, as an unlabeled packet, with a TTL that
       reflects the IP (network layer) forwarding.

If the incoming TTL is 1, only the first option applies.

In the multicast case, a meaningful LSP length or LSP segment length
is propagated to the FR-LSR egress node, enabling the egress to
decrement the TTL value before forwarding packets out of the non-TTL
LSP segment.

5.4.2 Performing MPLS TTL calculations

The calculation applied to the "input TTL" that yields the "output
TTL" depends on (i)the "input encapsulation", (ii)the "forwarding
encapsulation", and (iii)the "output encapsulation". The
relationship among (i),(ii), and (iii), can be defined as a function

   "D" of "input encapsulation" (ie), "forwarding encapsulation" (fe),
   and "output encapsulation" (oe). Subsequently the calculation
   applied to the "input TTL" to yield the "output TTL" can be described
   as:

      output TTL = input TTL - D(ie, fe, oe)

   or in a brief notation:

      output TTL = input TTL - d

   where "d" has three possible values: "0","1", or "the number of hops
   of the LSP segment":

   For unicast transmission:

d	Type of Input Encapsulation	Type of Forwarding Encapsulation	Type of Output Encapsulation
0	Frame Relay	Frame Relay	Frame Relay
1	any	Generic MPLS	Generic MPLS
number of hops of LSP segment	any	Generic MPLS or IP(network layer)	Frame Relay

   The "number of hops of the LSP segment" is the value of the "hop
   count" that is attached with the label used when the packet is
   forwarded, if LDP [LDP] has provided such a "hop count" value when it
   distributed the label for the LSP, that is the LDP message had a "hop
   count object". If LDP didn't provide a "hop count", or it provided
   an "unknown" value, the default value of the "number of hops of the
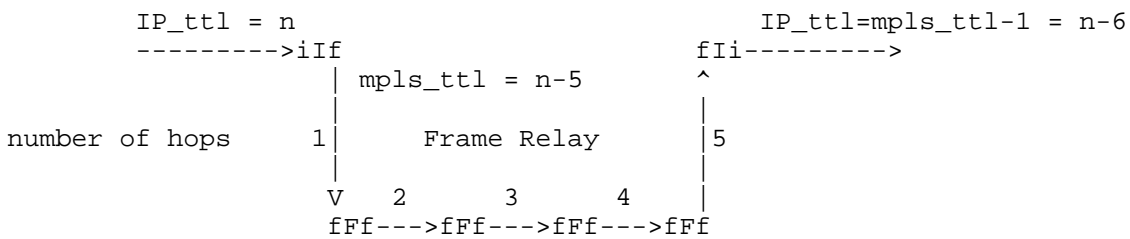   segment" is 1.

   When sending a label binding upstream, the "hop count" associated
   with the corresponding binding from downstream, if different than the
   "unknown" value, MUST be incremented by 1, and the result transmitted
   upstream as the hop count associated with the new binding (the
   "unknown" value is transmitted unchanged). If the new "hop count"
   value exceeds the "maximum" value, the FR-LSR MUST NOT pass the
   binding upstream, but instead MUST send an error upstream
   [LDP][ARCH].

For multicast transmission:

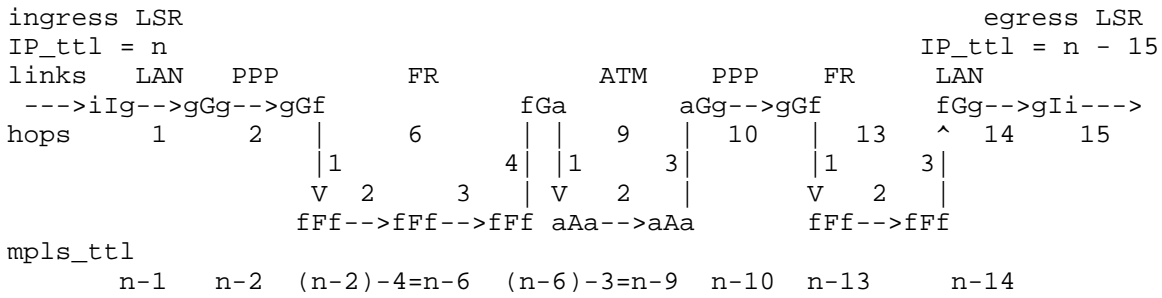| d | Type of Input Encapsulation | Type of Forwarding Encapsulation | Type of Output Encapsulation |
===	===	===	===
0	Frame Relay	Frame Relay	Frame Relay
1	any	Generic MPLS or IP(network layer)	Frame Relay
number of hops of LSP segment	Frame Relay	Generic MPLS or IP(network layer)	any

Referring to the "forwarding encapsulation" with the abbreviation "I"
for IP (network layer), "G" for Generic MPLS, and "F" for Frame Relay
MPLS, referring to an LSR interface with the abbreviation "i" if the
input or output encapsulation is IP and no MPLS encapsulation, "g"
when the input or output MPLS encapsulation is Generic MPLS, "f" when
it is Frame Relay, "a" when it is ATM, and furthermore considering
the symbols "iIf", "gGf", "fFf", etc... as LSRs with input,
forwarding and output encapsulations as referred above, the following
describes examples of TTL calculations for the Homogeneous and
Heterogeneous LSPs discussed in previous sections:

```
                    Homogeneous LSP
                    ---------------
        IP_ttl = n                            IP_ttl=mpls_ttl-1 = n-6
        --------->iIf                      fIi--------->
                  | mpls_ttl = n-5          ^
                  |                         |
number of hops   1|      Frame Relay        |5
                  |                         |
                  V   2      3      4       |
                  fFf--->fFf--->fFf--->fFf
```

 "iIf" is "ingress LSR" in Frame Relay LSP and
        calculates: mpls_ttl = IP_TTL - number of hops = n-5
 "fIi" is "egress LSR" from Frame Relay LSP, and
        calculates: IP_ttl = mpls_ttl-1 = n-6
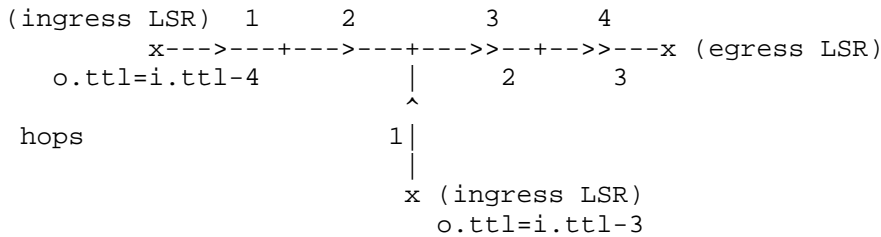
```
                        Heterogeneous LSP
                        -----------------
ingress LSR                                                egress LSR
IP_ttl = n                                               IP_ttl = n - 15
links  LAN   PPP         FR           ATM     PPP    FR    LAN
 --->iIg-->gGg-->gGf            fGa        aGg-->gGf      fGg-->gIi--->
hops   1     2   |     6      | |    9   | 10   | 13   ^  14     15
                 |1          4| |1     3|      |1     3|
                 V  2    3    | V   2   |      V   2    |
               fFf-->fFf-->fFf aAa-->aAa       fFf-->fFf
mpls_ttl
      n-1   n-2  (n-2)-4=n-6  (n-6)-3=n-9  n-10  n-13     n-14


"iIg" is "ingress LSR" in LSP; it calculates: mpls_ttl=n-1
"gGf" is "egress LSR" from Generic MPLS segment, and
      "ingress LSR" in Frame Relay segment and calculates: mpls_ttl=n-6
"fGa" "egress LSR" from Frame Relay segment, and
      "ingress LSR" in ATM segment and calculates: mpls_ttl=n-9
"gGf" is "egress LSR" from Generic MPLS segment, and
      "ingress LSR" in Frame Relay segment and calculates: mpls_ttl=n-13
"fGg" is "egress LSR" from Frame Relay segment, and
      ingress LSR" in Generic MPLS segment and calculates: mpls_ttl=n-14
"gIi" is "egress LSR" from  LSP and calculates: IP_ttl=n-15


      And further examples:

              Frame Relay Unicast -- TTL calculated at ingress

   (ingress LSR)  1      2        3       4
         x--->---+--->---+--->>--+-->>---x (egress LSR)
      o.ttl=i.ttl-4         |        2       3
                            ^
 hops                     1|
                           |
                          x (ingress LSR)
                            o.ttl=i.ttl-3


        Frame Relay Multicast -- TTL calculated at egress

            (egress LSR)x  o.ttl=i.ttl-3
   hops                 |
                        ^3
   (ingress LSR)        |              o.ttl=i.ttl-4
        x--->---+--->---+--->---+--->---x (egress LSR)
            1       2       3       4
```
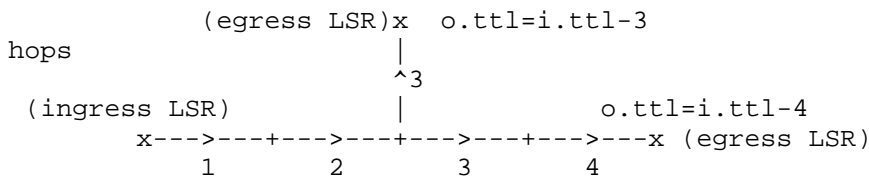
5.5 Label Processing by Ingress FR-LSRs

   When a packet first enters an MPLS domain, the packet is forwarded by
   normal network layer forwarding operations with the exception that
   the outgoing encapsulation will include an MPLS label stack [STACK]
   with at least one entry. The frame relay null encapsulation will
   carry information about the network layer protocol implicitly in the
   label, which MUST be associated only with that network protocol. The
   TTL field in the top label stack entry is filled with the network
   layer TTL (or hop limit) resulted after network layer forwarding
   [STACK]. The further FR-LSR processing is similar in both possible
   cases:

   (a) the LSP is homogeneous -- Frame Relay only -- and the FR-LSR is
   the ingress.

   (b) the LSP is heterogeneous -- Frame Relay, PPP, Ethernet, ATM,
   etc... segments form the LSP -- and the FR-LSR is the ingress into a
   Frame Relay segment.

   For unicast packets, the MPLS TTL SHOULD be decremented with the
   number of hops of the Frame Relay LSP (homogeneous), or Frame Relay
   segment of the LSP (heterogeneous). An LDP constructing the LSP
   SHOULD pass meaningful information to the ingress FR-LSR regarding
   the number of hops of the "non-TTL segment".

   For multicast packets, the MPLS TTL SHOULD be decremented by 1. An
   LDP constructing the LSP SHOULD pass meaningful information to the
   egress FR-LSR regarding the number of hops of the "non-TTL segment".

   Next, the MPLS encapsulated packet is passed down to the Frame Relay
   data link driver with the top label as output DLCI. The Frame Relay
   frame carrying the MPLS encapsulated packet is forwarded onto the
   Frame Relay VC to the next LSR.

5.6 Label Processing by Core FR-LSRs

   In a FR-LSR, the current (top) MPLS label is carried in the DLCI
   field of the Frame Relay data link layer header of the frame. Just
   as in conventional Frame Relay, for a frame arriving at an interface,
   the DLCI carried by the Frame Relay data link header is looked up in
   the DLCI Information Base, replaced with the correspondent output
   DLCI, and transmitted on the outgoing interface (forwarded to the
   next hop node).

The current label information is also carried in the top of the label
stack. In the top-level entry, all fields except the label
information, which is carried and switched in the Frame Relay frame
data link-layer header, are of current significance.

5.7 Label Processing by Egress FR-LSRs

When reaching the end of a Frame Relay LSP, the FR-LSR pops the label
stack [ARCH]. If the label popped is the last label, it is necessary
to determine the particular network layer protocol which is being
carried. The label stack carries no explicit information to identify
the network layer protocol. This must be inferred from the value of
the label which is popped from the stack.

If the label popped is not the last label, the previous top level
MPLS TTL is propagated to the new top label stack entry.

If the FR-LSR is the egress switch of a Frame Relay segment of a
hybrid LSP, and the end of the Frame Relay segment is not the end of
the LSP, the MPLS packet will be processed for forwarding onto the
next segment of the LSP based on the information held in the Next Hop
Label Forwarding Entry (NHLFE) [ARCH]. The output label is set to
the value from the NHLFE, and the MPLS TTL is decremented by the
appropriate value depending the type of the output interface and the
type of transmit operation (see section 6.3). Further, the MPLS
packet is forwarded according to the MPLS specifications for the
particular link of the next segment of the LSP.

For unicast packets, the MPLS TTL SHOULD be decremented by one if the
output interface is a generic one, or with the number of hops of the
next ATM segment of the LSP (heterogeneous), if the output interface
is an ATM (non-TTL) interface.

For multicast packets, the MPLS TTL SHOULD be decremented by the
number of hops of the FR segment being exited. An LDP constructing
the LSP SHOULD pass meaningful information to the egress FR-LSR
regarding the number of hops of the FR "non-TTL segment".

6. Label Switching Control Component for Frame Relay

To support label switching a Frame Relay Switch MUST implement the
control component of label switching, which consists primarily of
label allocation and maintenance procedures. Label binding
information MAY be communicated by several mechanisms, one of which
is the Label Distribution Protocol (LDP) [LDP].

Since the label switching control component uses information learned
directly from network layer routing protocols, this implies that the
switch MUST participate as a peer in these protocols (e.g., OSPF,
IS-IS).

In some cases, LSRs may use other protocols (e.g., RSVP, PIM, BGP) to
distribute label bindings. In these cases, a Frame Relay LSR should
participate in these protocols.

In the case where Frame Relay circuits are established via LDP, or
RSVP, or others, with no involvement from traditional Frame Relay
mechanisms, it is assumed that circuit establishing contractual
information such as input/output maximum frame size,
incoming/outgoing requested/agreed throughput, incoming/outgoing
acceptable throughput, incoming/outgoing burst size,
incoming/outgoing frame rate, used in transmitting, and congestion
control MAY be passed to the FR-LSRs through RSVP, or can be
statically configured. It is also assumed that congestion control
and frame header flagging as a consequence of congestion, would be
done by the FR-LSRs in a similar fashion as for traditional Frame
Relay circuits. With the goal of emulating a best-effort router as
default, the default VC parameters, in the absence of LDP, RSVP, or
other mechanisms participation to setting such parameters, should be
zero CIR, so that input policing will set the DE bit in incoming
frames, but no frames are dropped.

Control and state information for the circuits based on MPLS MAY be
communicated through LDP.

Support of label switching on a Frame Relay switch requires
conformance only to [FRF] (framing, bit-stuffing, headers, FCS)
except for section 2.3 (PVC control signaling procedures, aka LMI).
Q.933 signaling for PVCs and/or SVCs is not required. PVC and/or SVC
signaling may be used for non-MPLS (standard Frame Relay) PVCs and/or
SVCs when both are running on the same interface as MPLS, as
discussed in the next section.

6.1 Hybrid Switches (Ships in the Night)

The existence of the label switching control component on a Frame
Relay switch does not preclude the ability to support the Frame Relay
control component defined by the ITU and Frame Relay Forum on the
same switch and the same interfaces (NICs). The two control
components, label switching and those defined by ITU/Frame Relay
Forum, would operate independently.

Definition of how such a device operates is beyond the scope of this
document. However, only a small amount of information needs to be
consistent between the two control components, such as the portions
of the DLCI space which are available to each component.

7. Label Allocation and Maintenance Procedures

The mechanisms and message formats of a Label Distribution Protocol
are documented in [ARCH] and [LDP]. The "downstream-on-demand" label
allocation and maintenance mechanism discussed in this section MUST
be used by FR-LSRs that do not support VC merging, and it MAY also be
used by FR-LSRs that do support VC merging (note that this mechanism
applies to hop-by-hop routed traffic):

7.1 Edge LSR Behavior

Consider a member of the Edge Set of a FR-LSR domain. Assume that,
as a result of its routing calculations, it selects a FR-LSR as the
next hop of a certain route (FEC), and that the next hop is reachable
via a LC-Frame Relay interface. Assume that the next-hop FR-LSR is
an "LDP-peer" [ARCH][LDP]. The Edge LSR sends an LDP "request"
message for a label binding from the next hop, downstream LSR. When
the Edge LSR receives in response from the downstream LSR the label
binding information in an LDP "mapping" message, the label is stored
in the Label Information Base (LIB) as an outgoing label for that
FEC. The "mapping" message may contain the "hop count" object, which
represents the number of hops a packet will take to cross the FR-LSR
domain to the Egress FR-LSR when using this label. This information
may be stored for TTL calculation. Once this is done, the LSR may
use MPLS forwarding to transmit packets in that FEC.

When a member of the Edge Set of the FR-LSR domain receives an LDP
"request" message from a FR-LSR for a FEC, it means it is the
Egress-FR-LSR. It allocates a label, creates a new entry in its
Label Information Base (LIB), places that label in the incoming label
component of the entry, and returns (via LDP) a "mapping" message
containing the allocated label back upstream to the LDP peer that
originated the request. The "mapping" message contains the "hop
count" object value set to 1.

When a routing calculation causes an Edge LSR to change the next hop
for a route, and the former next hop was in the FR-LSR domain, the
Edge LSR should notify the former next hop (via an LDP "release"
message) that the label binding associated with the route is no
longer needed.

When a Frame Relay-LSR receives an LDP "request" message for a
certain route (FEC) from an LDP peer connected to the FR-LSR over a
LC-FR interface, the FR-LSR takes the following actions:

    - it allocates a label, creates a new entry in its Label
       Information Base (LIB), and places that label in the incoming
       label component of the entry;

    - it propagates the "request", by sending an LDP "request"
       message to the next hop LSR, downstream for that route (FEC);

In the "ordered control" mode [ARCH], the FR-LSR will wait for its
"request" to be responded from downstream with a "mapping" message
before returning the "mapping" upstream in response to a "request"
("ordered control" approach [ARCH]). In this case, the FR-LSR
increments the hop count it received from downstream and uses this
value in the "mapping" it returns upstream.

Alternatively, the FR-LSR may return the binding upstream without
waiting for a binding from downstream ("independent control" approach
[ARCH]). In this case, it uses a reserved value for hop count in the
"mapping", indicating that it is 'unknown'. The correct value for
hop count will be returned later, as described below.

Since both the "ordered" and "independent" control has advantages and
disadvantages, this is left as an implementation, or configuration
choice.

Once the FR-LSR receives in response the label binding in an LDP
"mapping" message from the next hop, it places the label into the
outgoing label component of the LIB entry.

Note that a FR-LSR, or a member of the edge set of a FR-LSR domain,
may receive multiple binding requests for the same route (FEC) from
the same FR-LSR. It must generate a new "mapping" for each "request"
(assuming adequate resources to do so), and retain any existing
mapping(s). For each "request" received, a FR-LSR should also
generate a new binding "request" toward the next hop for the route
(FEC).

When a routing calculation causes a FR-LSR to change the next hop for
a route (FEC), the FR-LSR should notify the former next hop (via an
LDP "release" message) that the label binding associated with the
route is no longer needed.

When a LSR receives a notification that a particular label binding is
no longer needed, the LSR may deallocate the label associated with
the binding, and destroy the binding. This mode is the "conservative

label retention mode" [ARCH]. In the case where a FR-LSR receives
such notification and destroys the binding, it should notify the next
hop for the route that the label binding is no longer needed. If a
LSR does not destroy the binding (the FR-LSR is configured in
"liberal label retention mode" [ARCH]), it may re-use the binding
only if it receives a request for the same route with the same hop
count as the request that originally caused the binding to be
created.

When a route changes, the label bindings are re-established from the
point where the route diverges from the previous route. LSRs
upstream of that point are (with one exception, noted below)
oblivious to the change. Whenever a LSR changes its next hop for a
particular route, if the new next hop is a FR-LSR or a member of the
edge set reachable via a LC-FR interface, then for each entry in its
LIB associated with the route the LSR should request (via LDP) a
binding from the new next hop.

When a FR-LSR receives a label binding from a downstream neighbor, it
may already have provided a corresponding label binding for this
route to an upstream neighbor, either because it is using
"independent control" or because the new binding from downstream is
the result of a routing change. In this case, it should extract the
hop count from the new binding and increment it by one. If the new
hop count is different from that which was previously conveyed to the
upstream neighbor (including the case where the upstream neighbor was
given the value 'unknown') the FR-LSR must notify the upstream
neighbor of the change. Each FR-LSR in turn increments the hop count
and passes it upstream until it reaches the ingress Edge LSR.

Whenever a FR-LSR originates a label binding request to its next hop
LSR as a result of receiving a label binding request from another
(upstream) LSR, and the request to the next hop LSR is not satisfied,
the FR-LSR should destroy the binding created in response to the
received request, and notify the requester (via an LDP "withdraw"
message).

When an LSR determines that it has lost its LDP session with another
LSR, the following actions are taken:

    - MUST discard any binding information learned via this
       connection;

    - For any label bindings that were created as a result of
       receiving label binding requests from the peer, the LSR may
       destroy these bindings (and deallocate labels associated with
       these binding).

7.2   Efficient use of label space - Merging FR-LSRs

   The above discussion assumes that an edge LSR will request one label
   for each prefix in its routing table that has a next hop in the FR-
   LSR domain. In fact, it is possible to significantly reduce the
   number of labels needed by having the edge LSR request instead one
   label for several routes. Use of many-to-one mappings between routes
   (address prefixes) and labels using the notion of Forwarding
   Equivalence Classes (as described in [ARCH]) provides a mechanism to
   conserve the number of labels.

   Note that conserving label space (VC merging) may be restricted in
   case the frame traffic requires Frame Relay fragmentation. The issue
   is that Frame Relay fragments must be transmitted in sequence, i.e.,
   fragments of distinct frames must not be interleaved. If the
   fragmenting FR-LSR ensures the transmission in sequence of all
   fragments of a frame, without interleaving with fragments of other
   frames, then label conservation (VC merging) can be performed.

   When label conservation is used, when a FR-LSR receives a binding
   request from an upstream LSR for a certain FEC, and it does already
   have an outgoing label binding for that FEC, it does not need to
   issue a downstream binding request. Instead, it may allocate an
   incoming label, and return that label in a binding to the upstream
   requester. Packets received from the requester, with that label as
   top label, will be forwarded after replacing the label with the
   existing outgoing label for that FEC. If the FR-LSR does not have an
   outgoing label binding for that FEC, but does have an outstanding
   request for one, it need not issue another request. This means that
   in a label conservation case, a FR-LSR must respond with a new
   binding for every upstream request, but it may need to send one
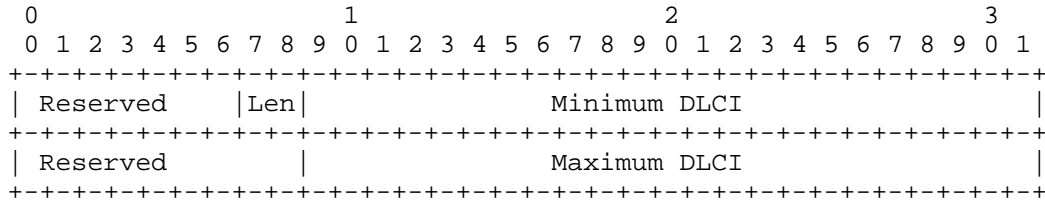   binding request downstream.

   In case of label conservation, if a change in the routing table
   causes FR-LSR to select a new next hop for one of its FECs, it MAY
   release the binding for that route from the former next hop. If it
   doesn't already have a corresponding binding for the new next hop, it
   must request one (note that the choice depends on the label retention
   mode [ARCH]).

   If a new binding is obtained, which contain a hop count that differs
   from that of the old binding, the FR-LSR must process the new hop
   count: increment by 1, if different than "unknown", and notify the
   upstream neighbors who have label bindings for this FEC of the new
   value. To ensure that loops will be detected, if the new hop count
   exceeds the "maximum" value, the label values for this FEC must be
   withdrawn from all upstream neighbors to whom a binding was
   previously sent.

7.3   LDP messages specific to Frame Relay

   The Label Distribution Protocol [LDP] messages exchanged between two
   Frame Relay "LDP-peer" LSRs may contain Frame Relay specific
   information such as:

   "Frame Relay Label Range":

```
       0                   1                   2                   3
       0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      | Reserved      |Len|               Minimum DLCI                |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      | Reserved          |               Maximum DLCI                |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

   with the following fields:

   Reserved
      This fields are reserved. They must be set to zero on
      transmission and must be ignored on receipt.

   Len
      This field specifies the number of bits of the DLCI. The
      following values are supported:

          Len   DLCI bits

          0   10
          2   23

   Len values 1 and 3 are reserved for future use.

   Minimum DLCI
      This 23 bit field is the binary value of the lower bound of a
      block of Data Link Connection Identifiers (DLCIs) that is
      supported by the originating FR-LSR. The Minimum DLCI should be
      right justified in this field and the preceding bits should be set
      to 0.

   Maximum DLCI
      This 23 bit field is the binary value of the upper bound of a
      block of Data Link Connection Identifiers (DLCIs) that is
      supported by the originating FR-LSR. The Maximum DLCI should be
      right justified in this field and the preceding bits should be set
      to 0.

"Frame Relay Merge":

```
        0 1 2 3 4 5 6 7
       +-+-+-+-+-+-+-+-+
       | Reserved     |M|
       +-+-+-+-+-+-+-+-+
```
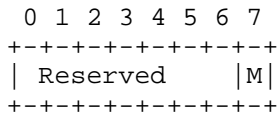
   with the following fields:

Merge
   One bit field that specifies the merge capabilities of the FR-LSR:
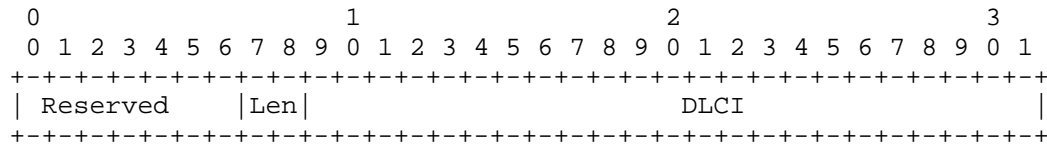
   Value   Meaning

    0   Merge NOT supported
    1   Merge supported

   A FR-LSR that supports VC merging MUST ensure that fragmented
   frames from distinct incoming DLCIs are not interleaved on the
   outgoing DLCI.

Reserved
   This field is reserved. It must be set to zero on transmission
   and must be ignored on receipt.

and "Frame Relay Label":

```
    0                   1                   2                   3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   | Reserved    |Len|                    DLCI                     |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

   with the following fields:

Reserved
   This field is reserved. It must be set to zero on transmission and
   must be ignored on receipt.

Len
   This field specifies the number of bits of the DLCI. The following
   values are supported:

       Len DLCI bits

        0   10
        2   23

Len values 1 and 3 are reserved for future use.

DLCI
   The binary value of the Frame Relay Label. The significant number
   of bits (10 or 23) of the label value are to be encoded into the
   Data Link Connection Identifier (DLCI) field when part of the
   Frame Relay data link header (see Section 4.).

8. Security Considerations

   This section looks at the security aspects of:

      (a) frame traffic,

      (b) label distribution.

   MPLS encapsulation has no effect on authenticated or encrypted
   network layer packets, that is IP packets that are authenticated or
   encrypted will incur no change.

   The MPLS protocol has no mechanisms of its own to protect against
   misdirection of packets or the impersonation of an LSR by accident or
   malicious intent.

   Altering by accident or forgery an existent label in the DLCI field
   of the Frame Relay data link layer header of a frame or one or more
   fields in a potentially following label stack affects the forwarding
   of that frame.

   The label distribution mechanism can be secured by applying the
   appropriate level of security to the underlying protocol carrying
   label information - authentication or encryption - see [LDP].

9. Acknowledgments

   The initial version of this document was derived from the Label
   Switching over ATM document [ATM].

   Thanks for the extensive reviewing and constructive comments from (in
   alphabetical order) Dan Harrington, Milan Merhar, Martin Mueller,
   Eric Rosen. Also thanks to George Swallow for the suggestion to use
   null encapsulation, and to Eric Gray for his reviewing.

   Also thanks to Nancy Feldman and Bob Thomas for their collaboration
   in including the LDP messages specific to Frame Relay LSRs.

10. References

   [MIFR]   Bradley, T., Brown, C. and A. Malis, "Multiprotocol
            Interconnect over Frame Relay", RFC 2427, September 1998.

   [ARCH]   Rosen, E., Callon, R. and A. Vishwanathan, "Multi-Protocol
            Label Switching Architecture", RFC 3031, January 2001.

   [LDP]   Andersson, L., Doolan, P., Feldman, N., Fredette, A. and R.
            Thomas, "Label Distribution Protocol", RFC 3036, January
            2001.

   [STACK] Rosen, E., Rehter, Y., Tappan, D., Farinacci, D., Fedorkow,
            G., Li, T. and A. Conta, "MPLS Label Stack Encoding", RFC
            3032, January 2001.

   [ATM]   Davie, B., Lawrence, J., McCloghrie, M., Rosen, E., Swallow,
            G., Rekhter, Y., and P. Doolan, "Use of Label Switching with
            ATM", RFC 3035, January 2001.

   [ITU]   International Telecommunications Union, "ISDN Data Link Layer
            Specification for Frame Mode Bearer Services", ITU-T
            Recommendation Q.922, 1992.

   [FRF]   Frame Relay Forum, User-to-Network Implementation Agreement
            (UNI), FRF 1.1, January 19, 1996.

11. Authors' Addresses

Alex Conta
Transwitch Corporation
3 Enterprise Drive
Shelton, CT 06484

Phone: 1-203-929-8810
EMail: aconta@txc.com

Paul Doolan
Ennovate Networks
60 Codman Hill Rd
Boxborough MA 01719

Phone: 1-978-263-2002
EMail: pdoolan@ennovatenetworks.com

Andrew G. Malis
Vivace Networks, Inc.
2730 Orchard Parkway
San Jose, CA 95134
USA

Phone: 1-408-383-7223
Fax: 1-408-904-4748
EMail: Andy.Malis@vivacenetworks.com

12. Full Copyright Statement

Acknowledgement