

Multilink Frame Relay

Prior to the advent of the multilink frame relay (MFR) protocol, some frame relay users were experiencing broadband application limits imposed by T1/E1 circuits. These users were approaching the T1/E1-T3/E3 bandwidth and price gap and were without an affordable, cost-effective option to acquire more bandwidth.

Consider the example of a headquarters office collaborating with its branch office. Using a T1/E1 network connection, the two offices could run a video conference application while users in each location download information from the Internet, send and receive email, or manage a user's forum on the company website. Eventually the bandwidth in the T1/E1 connection runs out and workflow and productivity are seriously impacted.

This example is typical of what is happening in enterprise networks today. Broadband applications, including transferring large files, streaming media, intranets, and extranets have become standard business practice. The fact is that more and more business is being conducted over networks. Bandwidth is in high demand as businesses move into real-time operations and adopt the Internet as a mission-critical tool. While the number of applications running over networks has grown, as has the number of users, bandwidth to the WAN has not kept pace.

Multilink frame relay was developed to address the need for more access bandwidth to support the ever-increasing number of broadband applications. MFR allows frame relay customers to realize cost-effective gains in bandwidth using inexpensive T1/E1 circuits without significant change to their network infrastructure.


Prior to MFR, service providers resorted to the use of inverse multiplexers to satisfy customers' needs for access bandwidth between T1/E1 and T3/E3. This, however, required the use of additional hardware at the customer site and at the carrier's central office.

What is MFR?

MFR is a packet transport protocol based on the adoption of the Frame Relay Forum's End-to-End MFR (FRF.15) and UNI/NNI MFR (FRF.16) Implementation Agreements. It provides physical interface emulation for frame relay devices by aggregating multiple physical circuits to transport data; however, the aggregation of physical circuits appears as one logical data pipe. While the principles of MFR can apply to bit rates such as NxDS0 or even NxT3, NxT1/E1 is garnering the major share of interest because of the cost or unavailability of high-speed circuits such as T3 and E3.

Bundling, a technique that combines physical interfaces and/or circuits into one or more logical data pipes, enables MFR access to support more total bandwidth than is available on any single physical interface. Bundling provides the benefits of the combined bandwidth of the aggregated physical interfaces, but with substantially less latency. Figure 1 shows the improved bandwidth and reduced latency that can be attained using a logical pipe consisting of bundled T1/E1 circuits transporting MFR fragments under the multilink frame relay protocol.

Figure 1: Single T1 vs. an MFR Bundle

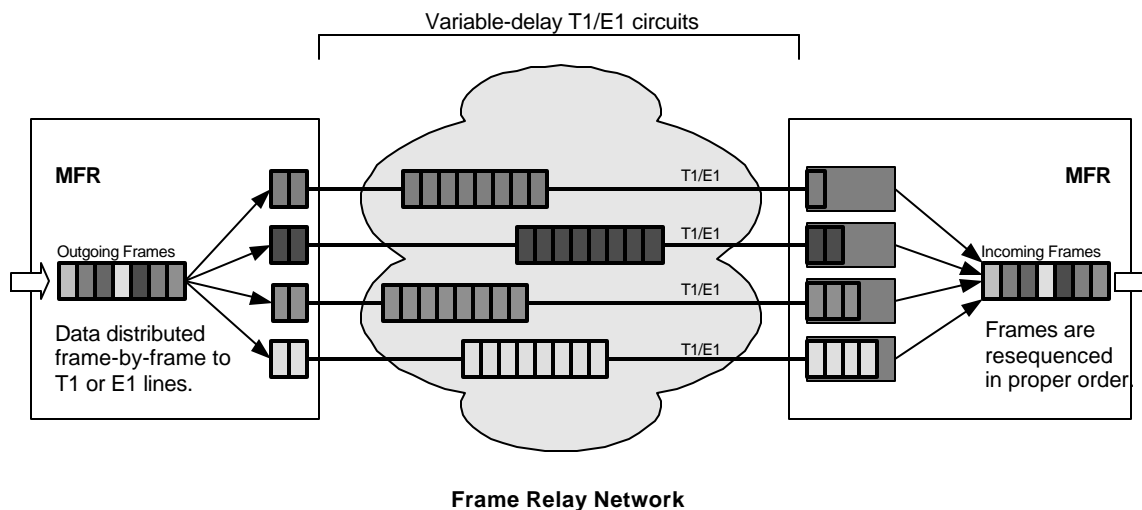


Greater service resilience is also provided because the bundle continues to function even when individual links within the bundle fail; the bundle automatically “throttles down” its bandwidth to the aggregate bandwidth of the remaining active links. In the example shown in Figure 1, the failure of one T1 circuit would result in the MFR bundle throttling down to 4.5 Mbps, thereby keeping the frame relay interface up and the data flowing. When the unavailable link comes back up, it is automatically added back to the bundle, restoring the original bandwidth of 6 Mbps.

UNI/NNI MFR (FRF.16)

Frame relay and MFR operate in two sub-layers within Layer 2 of the OSI model. In the “higher” sub-layer, a frame relay frame is transported over a permanent virtual circuit (PVC). In the “lower” sub-layer, the frame is distributed either in its entirety (Figure 2a) or in fragments (Figure 2b) across the T1 or E1 links at the Physical Layer (Layer 1). MFR frames or fragments are transported over a T1 or E1 link to the DCE at the Central Office (CO) where fragments are resequenced and reassembled in the correct order.

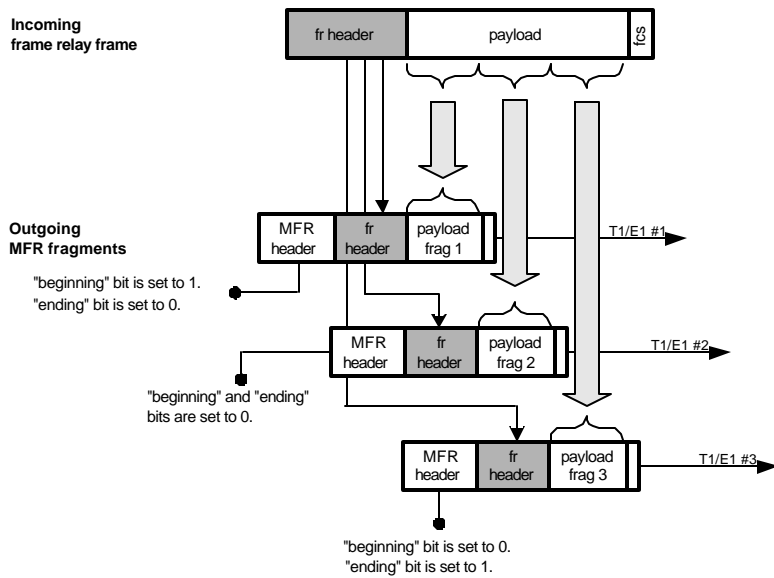
Figure 2a: Multilink Frame Relay UNI/NNI (FRF.16) Frame Transport



The sequence number in the MFR fragmentation header controls resequencing. Despite use of the fragmentation header to resequence frames sent across an MFR link, fragmentation is not required by the MFR specification. In many instances, algorithms deployed by MFR vendors optimize the transport of whole frames so that any delays that may occur are minimized. For example, fragmenting large frames will certainly lead to lower latency and increased throughput. However, smaller frames are transmitted as whole frame relay frames to avoid fragmentation/reassembly overhead. Should the option of fragmentation be deployed, however, both ends of the MFR link (the DTE and the DCE) must use the specifications for fragmentation used in the Frame Relay Forum FRF.12 Implementation Agreement.

Fragmentation tends to add a small amount of additional overhead to the MFR payload. In the example in Figure 2b, the incoming frame relay frame is broken into three MFR fragments. (User-configurable thresholds determine the number and the size of fragments.) The first fragment contains a copy of the original frame relay header. Each fragment contains an MFR header (used to reassemble the fragments in the correct order by the receiving endpoint), a duplicate of the original FR header (which preserves the characteristics of the original frame relay frame), and a payload fragment. The MFR header contains sequence numbers and beginning and ending bits. The first MFR header sets its beginning bit to one and ending bit to zero, and it establishes a sequence number for the related packets. Interim fragments set their beginning and ending bits to zero (to indicate an association with the data stream) and increment their sequence numbers. The last MFR header sets its beginning bit to zero and its ending bit to one (to indicate that it is the last fragment in the data stream), and it also increments its sequence number.

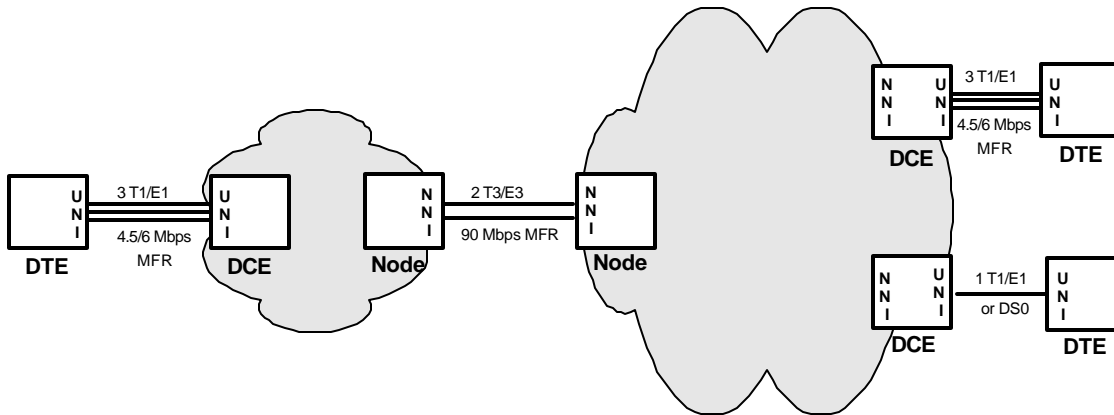
Figure 2b: UNI/NNI MFR Fragmentation (FRF.16)



The User-to-Network Interface (UNI) transports data between the DTE (data terminal equipment) and the central office DCE (connects users to the frame relay network) residing on the edge of the frame relay network. The Network-to-Network Interface (NNI) connects any two frame relay networks.

Figure 3 shows an example of an MFR application using both UNI and NNI interfaces. The DTE devices at remote user sites send MFR frames or encapsulated fragments to the DCE at the local Central Office (DCE) over copper T1 or E1 and NxT1/E1 lines (UNI/UNI). At the DCE, the MFR frames are correctly resequenced into their original order (if fragmentation is being used, fragments are reassembled into frame relay packets). The frames are then forwarded from the DCE to a switch in the network that has an MFR NNI to the second frame relay network. The frames are forwarded across the NNI to the second network and then to the access switch providing service to the CPE over T1/E1 or MFR T1/E1 circuits.

Figure 3: Multilink Frame Relay UNI/NNI, UNI/NNI, and NNI/NNI Connections



In terms of how they transport frames, UNI and NNI connections operate in exactly the same way. Differences between UNI and NNI connections occur in the LMI (Local Management Interface) layer, which will not be addressed in this document.

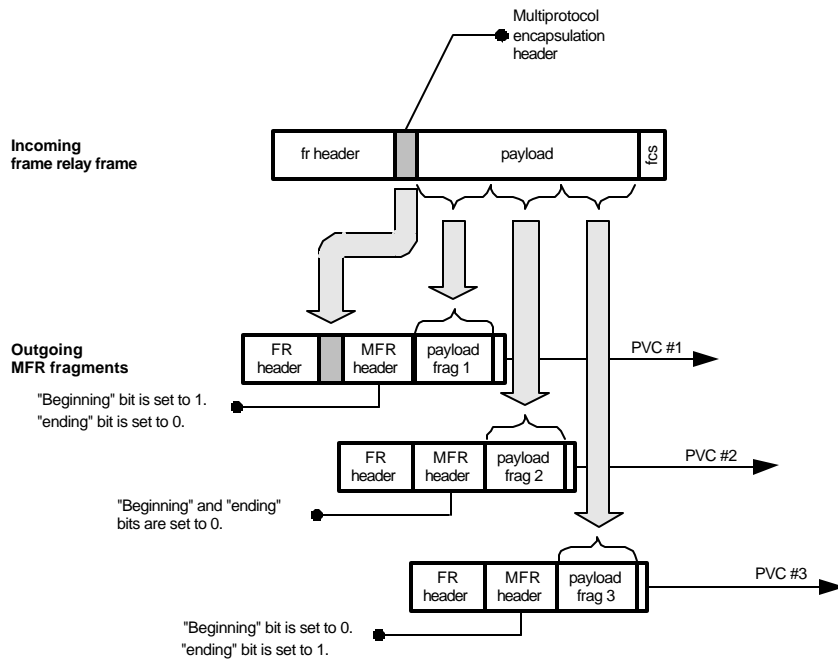
FRF.16 (UNI/NNI MFR) is implemented either between DTE and DCE or two directly connected nodes in two different networks. This is vastly different from FRF.15 (End-to-End MFR) implementation, which is configured between two DTE devices that can have any number of frame relay switches between them.

End-to-End MFR (FRF.15)

In cases where service providers do not support multilink frame relay (FRF.16), users requiring more bandwidth than can be provided by a single T1/E1 may opt for the End-to-End MFR implementation. Because existing frame switch equipment does not know how to handle MFR fragments, the MFR frames or fragments must be encapsulated into frame relay frames at the originating DTE and sent across the T1/E1 lines connected to the same or different service providers. Reassembly (if using fragments) and resequencing only occur at the destination DTE.

In an End-to-End (FRF.15) MFR network, MFR frames may also be sent in their entirety or in fragments. Figure 4 shows how a frame relay frame is fragmented for MFR processing in an End-to-End MFR circuit.

Figure 4: End-to-End MFR Fragmentation (FRF.15)



The incoming frame relay frame is broken into three MFR fragments, each containing a frame relay header, an MFR header (used to reassemble the fragments in the correct order by the receiving end), and the payload. The first fragment may contain the multiprotocol encapsulation header, which carries bits that provide traffic congestion notification across the network. These bits will be Ored from all the fragments at the time of reassembly to get the final congestion status.

The MFR header contains sequence numbers and beginning and ending bits. As in the UNI/NNI example above, the first MFR header sets its beginning bit to one and ending bit to zero, and it establishes a sequence number for the related packets. Interim fragments set their beginning and ending bits to zero (to indicate an association with the data stream) and increment their sequence numbers. The last MFR header sets its beginning bit to zero and its ending bit to one, and it also increments its sequence number. Intermediate FR switches only examine the 2-byte frame relay header and then transparently pass the MFR data on to the next network device. FRF 15 is transparent to the network.

Figure 5 shows an example of how encapsulation is used for End-to-End transmission. The original frame relay frame is fragmented and the subsequent MFR fragments are encapsulated into two new frame relay frames in preparation for transmission over the network.

Figure 5: End-to-End MFR Encapsulation (FRF.15)

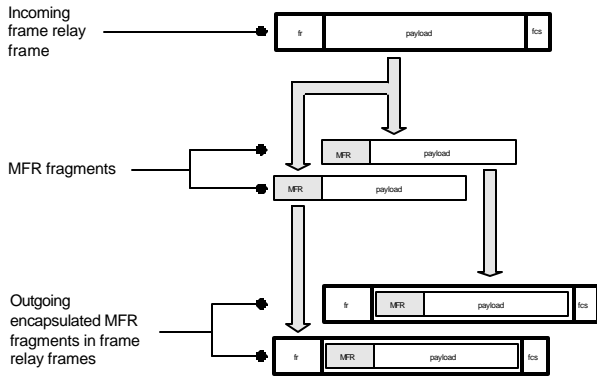
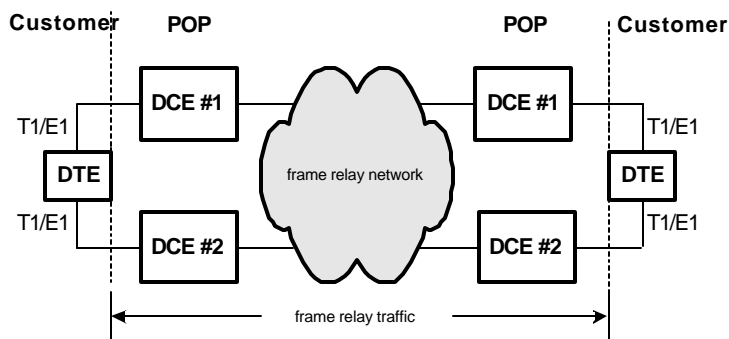


Figure 6, shows an example of End-to-End transmission. In this model, the originating DTE requires more bandwidth than a single T1 can provide and frame switches at point-of-presence (POP) locations cannot provide multilink support. Thus, the DTE must use multiple T1s (provided by the same service provider or two different service providers) to satisfy its bandwidth requirements.

Figure 6: End-to-End Transmission



The originating DTE encapsulates MFR frames or fragments within frame relay frames and then distributes this new frame traffic among the participating subscribers for transport over the frame relay network to the endpoint DTE. Frames may take varying paths enroute to the final destination and may arrive out of order. The DTE at the endpoint strips the frame relay headers and uses the sequencing information stored in the MFR headers to determine the correct order for reassembly and resequencing of the data.

Summary

The price gap between T1/E1 and T3/E3 service is substantial, and T1/E1 service is considerably more available than T3/E3 service. As T1/E1 pricing continues to fall, this price gap is widening. As this trend continues, T1/E1 becomes an even more viable option than T3/E3. The affordable and available characteristics of T1/E1, combined with the benefits of MFR, give service providers the ability to realize new revenue streams by providing flexible, low-cost network services that match individual customer requirements.

MFR also provides techniques that support the dynamic addition and removal of T1/E1 lines to change the total bandwidth available to the subscriber. A secondary result is the added resilience of multiple provisioned interfaces, such that when one interface or circuit fails, the emulated interface continues to support the frame relay service. A key value of this protocol is that in combining multiple physical interfaces or virtual circuits, a network operator can design a frame relay interface supporting more bandwidth than is available for any single physical interface or virtual circuit. And this can be accomplished with minimal change to the existing network infrastructure.

The introduction of frame relay marked a milestone in the data communications industry. It forever changed the way we think about data networks. The ongoing advancement and refinement of this protocol continues to allow frame relay to remain a prominent and popular data-traffic standard.