White Paper

August 8, 1999

www.wwgsolutions.com

Frame Relay Flow Control and Data Transmission Part 3: TFTP Over Frame Relay

Trivial File Transfer Protocol (TFTP) is a simple and basic method of reading and writing files. It provides IP-based file transfer, encapsulated in UDP. Since TFTP is encapsulated in UPD instead of TCP it must provide its own error control. Part 3 of this series discusses issues related to TFTP operation over Frame Relay networks. For a discussion of Frame Relay and TCP transport and flow control, see Parts 1 and 2 of this series, respectively.

TFTP Basics

Most TFTP servers and clients use the following standard set of parameters to transfer a file:

- File Transfer starts with either a Write Request or Read Request.
- Read and Write Requests contain the file name and sometimes include the file path.
- Segment numbering starts at 1 and is sequential.
- UDP encapsulation is standard on port 69.
- Each UDP encapsulated data segment contains 512 bytes of user data.
- The end of the file transfer is indicated by a frame with a valid checksum that contains less than 512 bytes of user data.
- The window size is set to 1, meaning that every segment must be acknowledged before the next segment is transmitted.
- The retransmission timer is generally set at 10 seconds, and is user adjustable in some instances.

TFTP was originally developed to work with BootP in diskless workstations. Because of TFTP's small size it could be contained within a single PROM. Over time other uses for TFTP were developed, such as network management platforms that use TFTP to load configurations and operating systems into routers.

Because only one frame is in transit at any time you might assume that TFTP is immune to congestion problems. In fact, congestion has an adverse affect on TFTP over Frame Relay. When using TFTP to transfer a file over Frame Relay to a remotely located device, the file transfer can appear to perform erratically. This poor performance typically manifests itself as excessively long file transfer times relative to the size of the file or, in more extreme cases, file transfer failure.

WAVETEK
WANDEL
GOLTERMANN
Communications Test Solutions

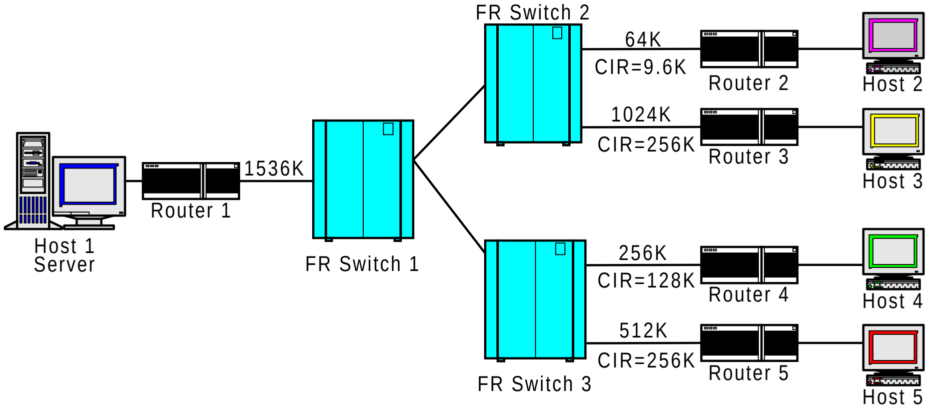

Figure 1

In Figure 1 the server is connected to the Frame Relay network through a full T1 interface. Host 2 is limited to 64 Kbit/s interface with a 9.6 Kbit/s CIR. When the network latency is low, multiple TFTP frames can be transmitted and acknowledged in less than one second. With each TFTP user data frame consuming about 4000 bits/s of bandwidth it only takes three frames to exceed the CIR. Depending upon how the network handles excess traffic, packets in excess of the CIR may be dropped.

A second source of file transfer problems occurs when the Frame Relay network drops TFTP packets or acknowledgments due to combined traffic levels. As discussed in Part 1 of this series, when the network is congested Frame Relay responds by dropping packets. The TFTP traffic load, combined with other network traffic, can trigger rate enforcement and result in TFTP packet loss.

In contrast to TCP, the TFTP protocol has a very rudimentary timer-based retransmission mechanism. The most common default setting for TFTP retransmission is 10 seconds. This means that a user data packet is retransmitted 10 seconds after it is sent if an acknowledgement is not received within that interval. This long retransmission pause can result in very slow file transfers. It is not uncommon for the long pauses in transmission to occur at regular, almost clock-like, intervals.

Consider this example: a host sends six packets and receives five acknowledgements, followed by a 10-second pause. At the end of the pause another six packets are sent, with the first packet being a retransmission of the unacknowledged packet, followed by another pause. This pattern continues until either the file transfer is complete or the TFTP server error stops the transfer. Once this pattern is established, even a relatively small file transfer can take an inordinate amount of time.

Unfortunately there is little that you can do to adjust for these problems. TFTP servers typically offer few configuration options that can be adjusted to improve performance. The only option in situations where the TFTP performance is unacceptable is to consider changing the Frame Relay network settings. In general, increasing the B_e helps prevent the switch from dropping packets, although the performance improvement may not justify the increased cost.

Because of the way TFTP and Frame Relay operate it is sometimes preferable to have relatively high network latency. Unfortunately there is no easy method of increasing the latency of a network. Attempting to use a router's traffic throughput enforcement capabilities only moves the problem from the network switch to the router. Until TFTP servers allow delay between the reception of an acknowledgment and the transmission of the next data frame, or allow the fixed retransmission timer setting to be decreased, this will continue to be a problem.

Some routers and FRADs offer bandwidth allocation mechanisms. If this is available for TFTP, exercise caution when choosing a value. Although a higher value may at first appear desirable, results tend to be better when less bandwidth is allocated. Remember that bandwidth allocation is a function of the router and the Frame Relay network; the server has no insight into these settings. The goal is to increase the packet latency over the network in order to slow the speed of the acknowledgments, thus reducing the number of packets sent per second. It will take some experimenting to determine the effect, if any, of various bandwidth allocation settings.

Identifying problems

In order to identify TFTP problems you need to capture some of the TFTP traffic with protocol analyzer. A WAN analyzer such as the DominoWAN or DominoHSSI analyzer is best, but you can also see the packet loss on a LAN segment using a LAN analyzer if that's what you have available. It isn't necessary to capture the entire session if you are experiencing file transfer failures or slow file transfers. Capturing a sample of the session should be sufficient to let you to determine if the TFTP problems are packet loss related.

Capture a sample of the TFTP traffic between the TFTP server and the client. After capturing the TFTP traffic, use a protocol decode engine such as WWG Examine™ to display a summary of the captured TFTP frames. Look at the following fields within the captured data to make a quick diagnosis of the TFTP problems:

- Use the timestamp field to find long pauses of about 10 seconds during which no TFTP packets are transmitted.
- Check for duplicate sequence numbers in data packets. Sequence numbers start at 1 and increment sequentially. Duplicate sequence numbers indicate retransmission.
- Check for missing sequence numbers, which indicate dropped packets.
- Look for a one-to-one relationship between data packets and acknowledgement packets. Any deviation indicates problems.

After examining the captured TFTP information you should have a good idea of the cause of the TFTP problems.

WAVETEK
WANDEL
GOLTERMANN
Communications Test Solutions