2114 West 7th Street
Tempe, AZ 85281 USA
Voice +1.480.333.2200
E-mail sales@comtechefdata.com
Web www.comtechefdata.com

Advantages of Using WAN Optimization with CDM-760 & ACM

January 2014

Satellite users have traditionally relied on worst-case link margin to overcome rain fade and other impairments, which can lead to significant inefficiencies. Adaptive Coding & Modulation (ACM) is a statistical, non-static advantage that enables dynamic changes in user throughput. ACM doesn't alter the power or the bandwidth of the carrier, it just changes the MODCODs so you can get more or less throughput according to link conditions. It allows the link to operate nearer the clear sky rate, throttling the capacity as environmental conditions dictate.

The benefits and value of ACM vary over time and are not guaranteed, but are predictable. ACM essentially converts link margin to an increase in the data throughput of satellite links and can increase overall link availability. With the ability to maximize throughput under all conditions (rain fade, inclined orbit satellite operation, interference and other impairments), ACM allows each remote to achieve maximum throughput, thereby maximizing network efficiency and availability.

When utilizing ACM in a satellite link configured with the CDM-760 Advanced High-Speed Trunking Modems, your data rate will change. The level of change depends on the link margin. But, the full range of the CDM-760 will take your link from 1 bit per symbol to 4.5 bits per symbol, which equates to an increase in data rate by > 4 times or a decrease by <1/4 times.

If your satellite link is fully utilized (filled pipe) and ACM causes your MODCOD to drop, the data rate of your link will be decreased and throughput will be reduced. If you are not utilizing WAN optimization tools to monitor the link and prioritize critical data, data will be randomly dropped and the overall throughput of the link will be severely impaired. This random discard of user packets will impact high value traffic and lower value traffic alike. This will also cause real-time based protocols to attempt re-transmissions and further congest the link.

WAN Optimization & Traffic Shaping

WAN optimization tools, such as our FX Series, monitor the traffic that comes onto to the network and enable you to perform traffic shaping with our multi-level Quality of Service (QoS). Traffic shaping involves delaying and sometimes dropping certain traffic on a link in order to optimize the link to a pre-defined profile.

There are basically two kinds of traffic on most networks.

- Reliable protocols – use some form of windowing and acknowledgments to control how much traffic is allowed to be "in the network". By selectively delaying some of the traffic, the natural response of the connection is to respond by slowing the rate of that connection.
- Real-time protocols – typically send data at a fixed data rate. Some of this data is also very sensitive to latency. Typically this latency-sensitive traffic, if it is delayed long enough, might as well be dropped.

Traffic shaping is a multi-step process that involves monitoring the WAN capacity, classifying the traffic and separating it into various queues, followed by draining the queues according to the appropriate drain algorithms. These steps are detailed below.

Real-time protocols need to be given the proper priority in order to get the appropriate access to the network as capacity diminishes. Reliable protocols will eventually adjust their data rate to the actual available capacity, but the process of adaptation itself wastes bandwidth and is not an efficient or even necessarily a fair process. Traffic shaping with QoS ensures that the traffic the network operator wants on the link gets the link.

Monitor WAN Capacity

When the WAN optimization is deployed in front of the modem as depicted below, it constantly monitors the transmit data rate of the modem and the WAN capacity. It knows how much congestion the link is experiencing. If the data rate is constricted due to changing MODCODs and the committed error rate cannot be met, the FX Series can apply the QoS rules to ensure that the lower priority (value) packets are dropped first.

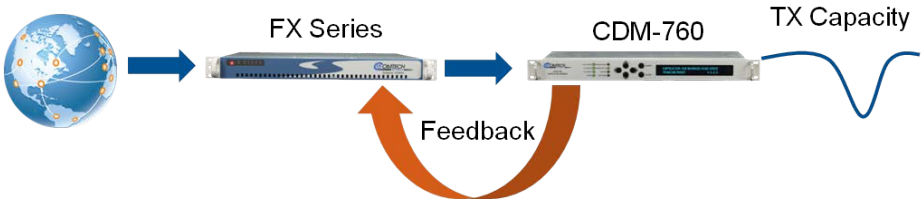

Ingress Data

The data traffic destined for the WAN enters the FX Series.

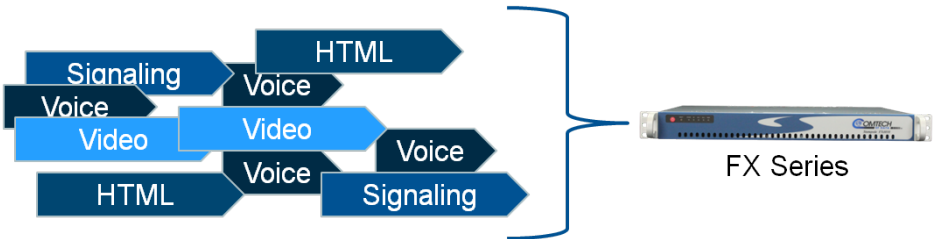

Prioritize Data Based on Filter Rules

The data traffic destined for the WAN is labeled and sorted by the FX Series. Various classifications and filters can be utilized to assign priorities on a per packet basis. Each packet is directed into a queue, where it can be handled by the second phase of the traffic shaping engine. Traffic classification can be done by:

- VLAN
- Source or destination IP address
- Source or destination port
- Protocol
- DSCP bits/ToS

Filters are defined using these fields, or combinations of these fields. Each filter directs traffic into a queue. More than one filter can direct traffic into the same queue. One of the issues in shaping is that traffic can match more than one filter. In the FX Series, filters are prioritized; the first filter that is matched is the filter that will be used.

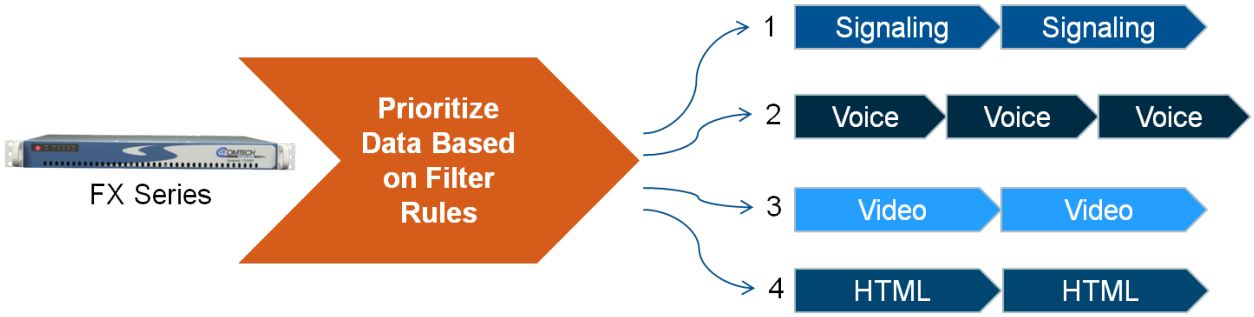

Drain Data Based on Rules

Once the traffic has been filtered and put in different priority classifications, it is drained based on established rules. Two basic drain algorithms are supported:

- Strict Priority – drains queues in order of priority. Packets in higher priority queues will be sent before lower priority packets. Queues of the same priority will be handled in fair manner; traffic flows in a manner proportional to the total traffic in the queues. As the highest priority queues suffer little or no latency, traffic that is latency-sensitive should be directed into the higher priority queues.
- Min-Max – Besides priority, there are two additional parameters, committed information rate (CIR) and maximum information rate (MIR). If possible, the requested CIR for each queue will be met. If the requested CIR cannot be met, then traffic will be dropped, starting with the lowest priority queues. The queues are first delayed, and then packets are dropped. When more traffic can be supported than the requested CIR, each queue is allowed to drain traffic up to the MIR, again in priority order.

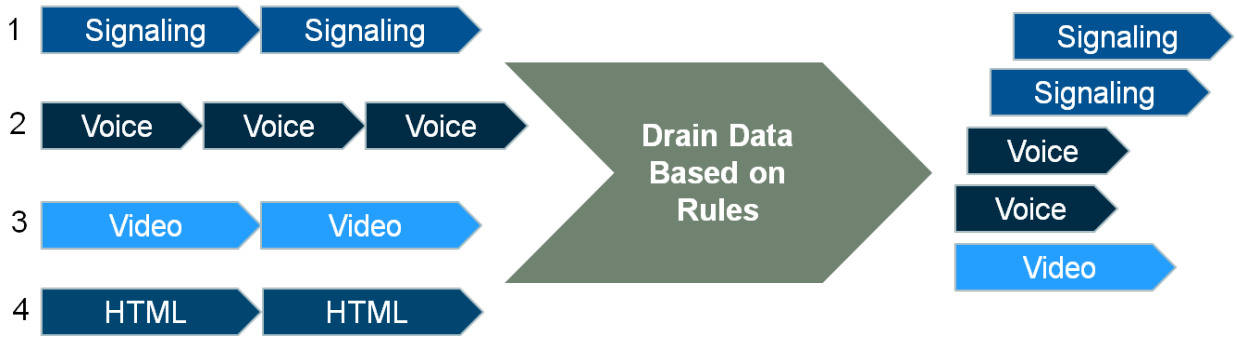

Summary

When the CDM-760 is operating in ACM mode in conjunction with the FX Series, the FX Series will update its data rate multiple times a second. The drain algorithms are followed for each of the queues, but the total traffic will match the link rate. As an ACM link gets impaired, the data rate will start to drop. If the link is full, this will, by necessity, require that some packets get dropped until the overall traffic gets in synch with the link rate.

This combination of technologies ensures the highest service quality with minimal jitter and latency for real-time traffic. It facilitates priority treatment of mission-critical applications by allowing the link to process critical traffic while dropping non-critical traffic during low-bandwidth conditions. It enables real-time traffic and other low priority traffic to seamlessly co-exist on the same link without impacting voice quality or mission-critical data delivery.

The combination of the CDM-760 and WAN optimization is not always required. In the unlikely scenario where you are just transmitting traffic with the same “value,” you don’t necessarily need the optimization provided by QoS and traffic shaping. However, virtually all large circuits where the CDM-760 Advanced High-Speed Trunking Modem would be utilized have a mix of various traffic types that are transported over satellite links. Some of the traffic will have more value and some of the traffic will have less value. This is where the full benefits of the CDM-760 + WAN optimization can be realized.

For additional information, please contact us.

Email: sales@comtechefdata.com
Voice: +1.480.333.2200
Web: www.comtechefdata.com