



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

Bridged Point-to-Multipoint: Layer 2 Connectivity in a Hub-Spoke Satellite Network

A technical description of the implementation of hub-spoke networks with transparent Layer 2 bridged Ethernet connectivity using SLM-5650A Satellite Modems.

June 2011



Abstract

Satellite networks are often designed to support connectivity for Internet Protocol (IP) data traffic. In many networks (particularly Government networks), data traffic is encrypted prior to arriving at the satellite communications element of the network. In encrypted IP-based networks, it is often highly desirable to have the satellite network transparently bridge traffic (i.e. operate at “Layer 2” in the OSI model, rather than Layer 3 or higher). This eliminates the requirement to support routing protocols and other Layer 3 functions in the satellite communication equipment on the “black” side of the encryptor, which in turn greatly simplifies the configuration and operation of the overall network.

In many IP-based satellite networks, a point-to-multipoint, or “hub-spoke” network architecture is also desired. However, standard Layer 2 Ethernet switches or other networking products operating at Layer 2 do not support hub-spoke networks. Hence without special processing, the requirement for Layer 2 connectivity in the satellite network conflicts with the requirement for a point-to-multipoint network architecture.

Comtech EF Data’s bridged point-to-multipoint implementation provides a solution to this problem, and enables Layer 2 bridged connectivity in a point-to-multipoint network.

Point-to-Multipoint Network Architecture

In a hub-spoke architecture, distributed remote terminals communicate to a central hub location. “Point-to-multipoint” is another term that is used to describe this architecture.

Figure 1 shows a hub-spoke architecture connecting a number of secure Government networks using SLM-5650A Satellite Modems. In this network, data traffic from the hub to the remotes is time division multiplexed (TDM) into a common Forward Link (FL). All remote terminals receive the FL. Remote-to-hub traffic is transmitted by individual Return Links (RLs) from each remote.

In networks where multiple remotes need to connect to a common hub, a network architecture using a shared FL is often preferred because satellite bandwidth can be utilized more efficiently than an alternative architecture using multiple point-to-point connections to each remote. This is particularly true when the network is IP-based, with time varying data rates to the remote users. Key reasons for this efficiency advantage include:

(a) Broadcast (Multicast) Traffic:

Broadcast packets destined to all remotes are transmitted only once in the shared FL. An alternative architecture using multiple point-to-point links would need to transmit the broadcast packets in each link.

(b) Statistical Multiplexing:

Shared FL capacity is utilized by the remote terminals that are active at a given time (i.e. capacity is not dedicated to idle terminals, nor is excess capacity dedicated to low data rate terminals). If data rates of the remotes vary over time, the ability to share the FL capacity provides both higher peak throughput and higher average throughputs to the remotes.

Layer 2 connectivity is enabled by the “BPM Network” function shown in Figure 1. This function enables the satellite network to appear as a bridged LAN from the perspective of the Government networks on the red-side of the encryptors (as shown in the bottom part of Figure 1).

The BPM network function is implemented by a combination of packet processing in the SLM-5650A, and configuration in external managed switches. Details are provided in Section 0.

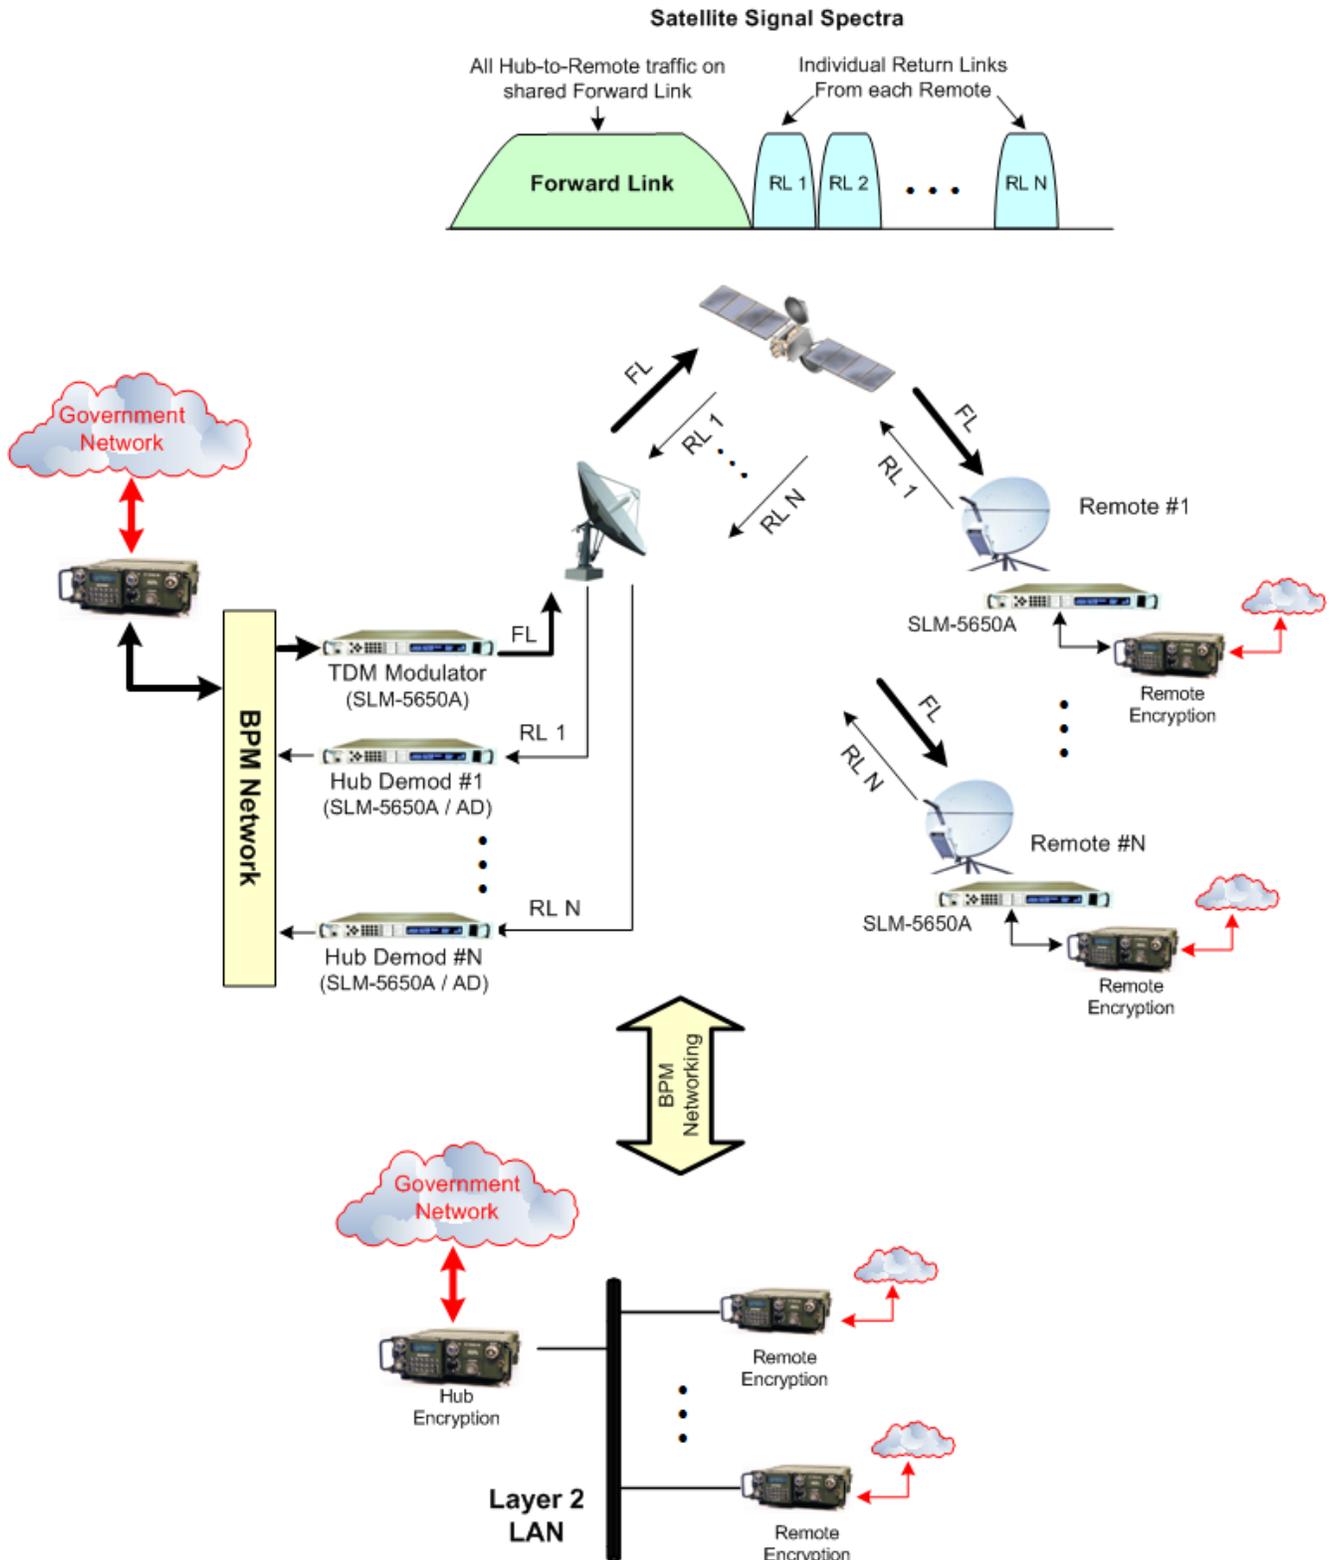


Figure 1: Point-to-Multipoint Satellite Network Provides Layer 2 Connectivity for Secure Government Networks

Bridged Point-to-Multipoint Implementation

A block diagram showing the key elements of the BPM solution at the hub is shown in Figure 2.

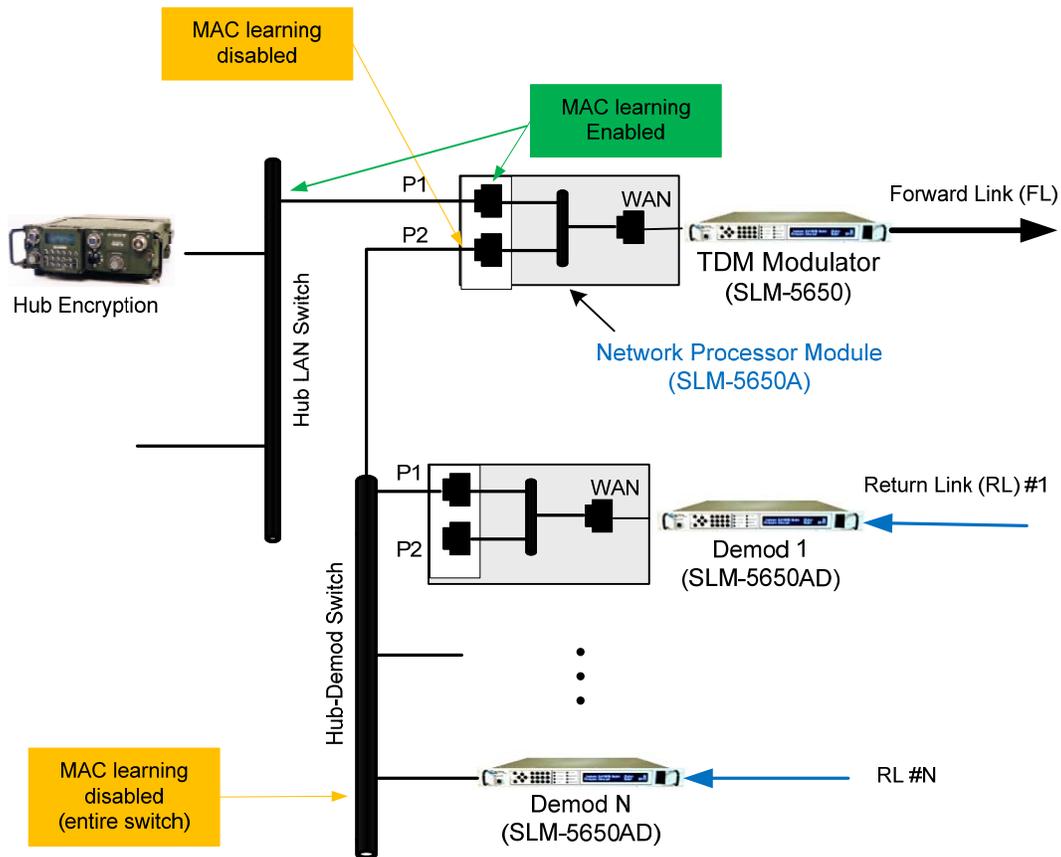


Figure 2: Bridged Point-to-Multipoint Block Diagram

Due to the split-path topology at the hub, the traffic in the forward and return link directions is processed by different SLM-5650A or SLM-5650AD (demod only) devices. In the forward link direction, all traffic is transmitted by the hub TDM modulator (a SLM-5650A acting as a shared modulator). In the return link direction, packets received via single carrier per channel (SCPC) return links are processed by hub demods (SLM-5650ADs operating as receive-only devices).

Both the hub TDM modulator and the demods are configured with Network Processor modules, which have 4-port Ethernet switch interfaces (2 of the 4 ports of the Network Processor module are shown in Figure 2). The external hub Ethernet switch supporting data traffic is shown as two separate switches, the “Hub LAN Switch”, and the “Hub Demod Switch”. It is important to note that this is a functional concept only. The physical implementation of this switch could be accomplished in at least one of three ways¹:

- a. Two separate switches, as illustrated in Figure 2
- b. One physical switch, partitioned into two logical switches using port isolation (many commercially available managed switches have this capability)
- c. Hub LAN switch implemented as a single external switch, with the functionality of the hub-demod switch implemented by means of daisy-chaining the LAN ports of the hub demod units together (i.e. connecting the P1 port of one hub demod to P2 of the next hub demod). This has the advantage of reducing the required hub equipment.

In the forward link direction, packets received from the hub data network ingress to the TDM modulator on a LAN port (P1), are forwarded to the satellite WAN port, encapsulated, optionally encrypted, and then transmitted to the satellite. No packets

¹ Note that if there is only one device (e.g. the hub encryption device) connected to the Hub LAN switch, then this switch element is unnecessary, and the hub encryption device would connect directly to P1 of the TDM modulator. Use of the daisy-chaining mechanism of implementing the hub demod switch function could then eliminate the need for any external hub switch equipment.

need ever be forwarded from the hub TDM modulator to the hub demods (hub demods are receive-only from the satellite direction). Hence a filter rule is put in place to block any data traffic input to the hub TDM port connected to the hub LAN switch (P1) being output from the port connected to the hub demod switch (P2).

For return link packets arriving via SCPC channels, care must be taken to avoid potential issues associated with the split path topology inherent in the point-to-multipoint system architecture. Two issues need to be addressed:

1. MAC address learning in the hub switches and modem switch ports
2. MAC address filtering to avoid unwanted packet transmission and/or reception

MAC Learning in the Hub Switches:

Typically, a Layer 2 switch learns the MAC-to-port association (MAC learning) during operation. However, in managed switches, MAC Learning can optionally be disabled on a per port basis. Disabling MAC learning is necessary in two places:

- (1) LAN port of the hub TDM modulator Network Processor that is connected to the hub demod switch (port P2)
- (2) External hub demod switch ports that are connected to each of the hub demods (all ports).

Disabling MAC learning on these ports avoids confusion in the hub layer 2 switches². In the return link direction, disabling MAC learning on the hub demod switch ports causes all return link packets to be broadcast to all ports on the switch, including the port connected to the LAN port of the hub TDM modulator (P2). This serves to aggregate all of the return link traffic from the remotes sites into hub TDM modulator port P2.

The hub TDM modulator P1 port and WAN port are configured with MAC learning on. The hub TDM modulator P2 port is configured with MAC learning off. In addition, a feature of the TDM modulator switch called port association is enabled. This feature allows the source MAC addresses from the TDM modulator P2 port to populate the MAC learning table for the TDM modulator WAN port.

As a result, when a return link packet arrives at the TDM modulator port P2, based upon the learned MAC address, the packet is sent to the correct LAN or WAN destination. That is, the packet is sent to the hub TDM modulator P1 port if it came from a remote and is destined for a device on the hub network. If the packet came from a remote and is destined to another remote network, the packet is sent to the hub TDM WAN port for transmission to the forward link WAN TDM carrier.

Shared Outbound MAC Filtering at the Remote Modems

The hub TDM modulator will broadcast hub to remote and remote to remote packets as appropriate on the forward link. Because the outbound TDM carrier is shared across all remotes, every remote receives a copy of each packet. For large shared forward links, the amount of data traffic can sometimes exceed the ability of the subsequent Layer 3 devices at the remote to process, and hence a filtering rule to extract only packets destined for a given remote is optionally implemented.

All unicast Ethernet data packets destined for a given remote terminal will have an Ethernet Destination Address (DA) of the router or encryption device connected to the remote modem. MAC filtering allows only unicast packets with a DA matching this device, plus all multicast (broadcast) packets to be forwarded. The unicast DA filter address is an operator configurable parameter

Conclusion

Bridged point-to-multipoint enables a satellite operator to implement a point-to-multipoint network that transparently bridges Ethernet traffic. This feature enables the operator to combine the bandwidth efficiency benefits of a point-to-multipoint network topology, with the simplicity of a bridged satellite network. Contact us for more information.

Email: sales@comtechefdata.com
Voice: 480.333.2200
Fax: 480.333.2540
Web: www.comtechefdata.com

² For example, consider the case where a given Remote with Source Address = "A", sends a packet to a device on the Hub network. If MAC Learning where enabled on port "P2" shown the Figure, then the MAC-to-port association would map MAC Address "A" to port P2. If a packet generated at the Hub destined for Remote Address A arrived at port P1, the switch would send this packet to port P2, rather than the intended satellite WAN port. By disabling MAC learning on P2, the packet is transmitted as desired to the satellite WAN port.