

# ETHERNET PACKET FILTERING for FTI – PART1

Ø. Holmeide<sup>1</sup>, J-F. Gauvin<sup>2</sup>

1 OnTime Networks AS, Oslo, Norway  
oeyvind@ontimenet.com

2 OnTime Networks AS, Oslo, Norway  
Jean.frederic@ontimenet.com

## Abstract:

Today's modern flight test systems are based to a large extent on Ethernet technologies. The never ending demand for increased bandwidth and speed also rises the need for clever packet filtering solutions on the Ethernet switches. The challenge is to avoid network bottlenecks and to ensure that Ethernet end nodes are not overwhelmed with data not meant for the given node. This paper describes the different packet filtering methods and how these techniques can be optimized for Flight Test Instrumentation (FTI) use.

**Keywords:** VLAN, IP, UDP, multicast, broadcast, packet filtering

## Introduction

The introduction of Ethernet in Flight Test Instrumentation (FTI) networks resulted in a high bandwidth communication solution that outperforms other legacy communication technologies by far. First generation FTI networks were based on 100Mbps, while the majority of today's networks are based on gigabit. Future FTI networks will most likely be based on switches with both gigabit and 10gigabit ports. This high bandwidth is great, but there is one concern: standard Ethernet is not deterministic. An Ethernet switch or a recorder may suffer from congestions unless the network is properly engineered.

The Ethernet switch provides the communication backbone for the FTI network and is a core element in order to ensure that packets sent between two end nodes are not lost. A fully manageable switch can filter packets, and avoid packet loss due to congestion or unnecessary packet filtering at a recorder received. This paper describes packet filtering based on Virtual Local Area Network (VLAN) (layer 2), Internet Protocol (IP) (layer 3) and User Datagram Protocol (UDP) (layer 4).

## OSI model

The seven layer model known as the Open Systems Interconnect (OSI) model is a standard established by the International Organization for Standardization (ISO) committee.




Fig. 1: The 7 layers of the OSI model-

The seven layers of the OSI model (ref Fig. 1) are as follows

- Layer 1, physical layer
- Layer 2, data link layer
- Layer 3, network layer
- Layer 4, transport layer
- Layer 5, session layer
- Layer6, presentation layer
- Layer 7, application layer

This paper will cover the packet filtering based on layer 2, 3 and 4.

The physical layer is a direct point-to-point data connection, commonly known as PHY, and consists of the hardware networking technology.

The data link layer is a reliable direct point-to-point data connection, commonly known as

MAC, this layer provides a reliable link between two directly connected nodes by detecting errors that may occur in the physical layer

The network layer addresses routing and delivery of datagrams between points on a network, commonly known as IP, this layer provides the functional and procedural means of transferring variable-length data sequences from a source to one or several destination host.

The transport layer handles the delivery of packets between points on a network, for FTI, this is commonly known as UDP (also TCP, DCCP, SCTP can be a part of this layer), this layer provides end-to-end communication services for applications.

Packet filtering can be based on each of the above layers depending on the complexity and the network applications that are used by the end nodes.

**VLAN packet filtering**

MAC packet filtering can be performed by using the VLAN (Virtual Local Area Networks) configuration. The principle of VLAN packet tagging was introduced in the Institute of Electrical and Electronics Engineers (IEEE) 802.1Q standard [2]. VLAN are used to segregate traffic in a network. This technique limits broadcast, multicast and flooded unicast network traffic by introducing VLAN segments. VLAN can be used for both data and management traffic.

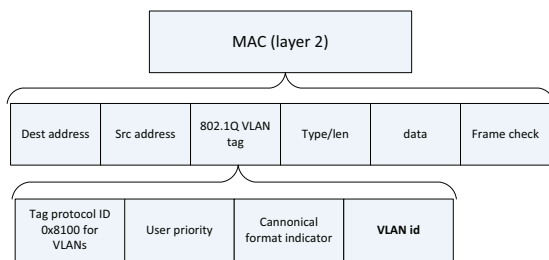


Fig. 2: IEEE802.1Q VLAN tag description

VLAN tagged packets can be used on all links, but normally not on the switch edge ports, where the end nodes are connected. The switch-to-switch ports (the trunk ports) can be configured to send packets with VLAN tags. A switch port that supports VLAN tags can be member of several VLANs. This allows overlapping VLANs. The switches must support the IEEE802.1Q [2] protocol.

The four bytes VLAN tag (ref Fig. 2) is assigned to the Ethernet header. Up to 4096 VLANs can be configured on an IEEE802.1Q compliant switch.

VLAN settings can be scaled from a simple application containing one single switch up to a large amount of network elements, such as switches and routers.

VLAN tags are inserted and removed by the switches in a network setup where the switches support IEEE802.1Q and the end nodes do not. The default VLAN ID of a given switch port is inserted by the switch upon arrival of an untagged packet on the switch port. The VLAN tag is removed when the packet is forwarded on a port connected to an end node. A tagged VLAN packet received on a switch port will be dropped if the VLAN ID is not configured on the given port.

**Port based VLAN – static VLAN**




Fig. 3: Simple port based VLAN application

The example shown in Fig. 3 demonstrates a setup using a single switch with port based VLAN as the VLAN technique. The VLAN is used to segregate the network traffic. VLAN 10 is configured to allow communication between IPCamera1, IPCamera2 and Recorder2 and VLAN 20 is used for communication between Data Acquisition Unit (DAU1), DAU2 and Recorder 1. The PC, used for monitoring the whole network, is a member of both VLANs and can send or receive data to/from any network element while the elements that belong to VLAN 10 cannot communicate with any of the network elements belonging to VLAN 20 and vice versa.

Configuring VLAN at switch port level means to assign a switch port to a certain VLAN ID. An untagged packet will be sent to the switch, the switch core will insert a 802.1Q [2] tag in the MAC header based on the port VLAN configuration, and then the packet will be forwarded to one or more ports belonging to the same VLAN as the receiving port. The VLAN tag will be removed at the egress port(s).

This type of simple VLAN setup allows a segregation of network segments. Broadcast, multicast or flooded unicast traffic belonging to a given VLAN will be kept in this VLAN. In the previous example no traffic from VLAN 10 will interfere with VLAN 20 traffic.

### Dynamic VLAN

Large networks may encounter the need to have multiple VLAN running on the same physical link, i.e. overlapping VLANs. In an 802.1Q [2] enabled network consisting of at least 2 switches, where multiple VLAN needs to coexist on a single link, TRUNK port(s) must be defined.

An ACCESS port is a port on the edge of the network. This port is in most cases connected to an end node such as a Data Acquisition Unit, recorder, PC, etc. The VLAN number assigned to an ACCESS port, i.e. default VLAN ID, corresponds to the VLAN on which the end node is allowed to communicate. The switch will insert a VLAN tag on all incoming packets on a switch port defined as an ACCESS port and remove the VLAN tag on all packets egressing the port.

A switch port defined as a TRUNK port is a port that is connected to another IEEE802.1Q

capable network device (e.g. another switch or a router). This port is used as a bundle for all the VLANs that need to coexist on this physical link. A TRUNK port can handle tagged packets.

The switch ports configured as TRUNK ports automatically propagate VLAN information from network device to network device. This can be done by the Generic Attribute Registration Protocol (GARP) VLAN Registration Protocol (GVRP). With GVRP, a single switch is configured with the desired end node VLANs on the ACCESS ports and other GVRP enabled switches in the network will learn those VLANs dynamically. An end node can be plugged anywhere in the network and then be connected to all other end nodes belonging to the same VLAN.

Let us extend the network setup described in Fig. 3, with the use of several GVRP enabled switches, see Fig. 4. VLAN 10 and VLAN 20 will need to share the links between switches 1 and 2 and between switches 2 and 3. These switch ports are configured as TRUNK ports allowing VLAN tagged packets with VLAN 10 and 20 to be forwarded between the switches. Each end node is connected to an ACCESS port configured with VLAN 10 or 20.

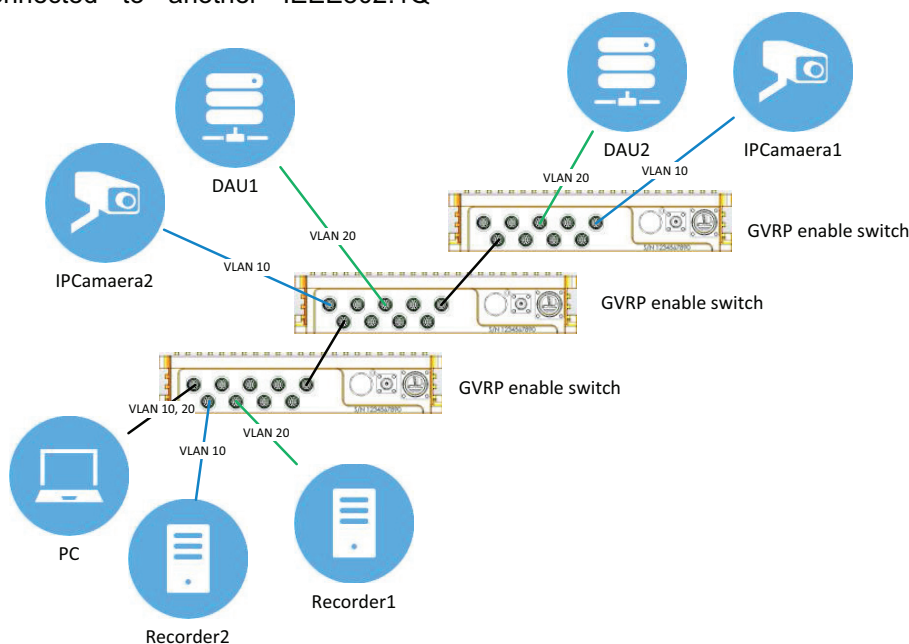


Fig. 4: VLAN using GVRP application

### IP packet filtering

Packet filtering based on IP destination addresses can be used in order to reduce the network load on the network links. The idea is to only forward those IP packets on a given link that is required in order to avoid network

congestion on the switches or exceeding the recorder storage capacity.

An IP packet can be sent with a unicast, multicast or broadcast destination address. This chapter focuses on IP packet filtering based on multicast IP destination addresses. IP multicast packet filtering is used widely in FTI systems. This technique allows one multicast

producer to send data to several multicast consumers.

The following IP range is defined as multicast IP addresses: [224.0.0.1 .. 239.255.255.255].

Let us consider the following application, where two data acquisition units (DAU1 and DAU2) are sending two multicast streams each to the network. Each multicast stream is designated to Recorder2 in addition to one more multicast consumer in the network.

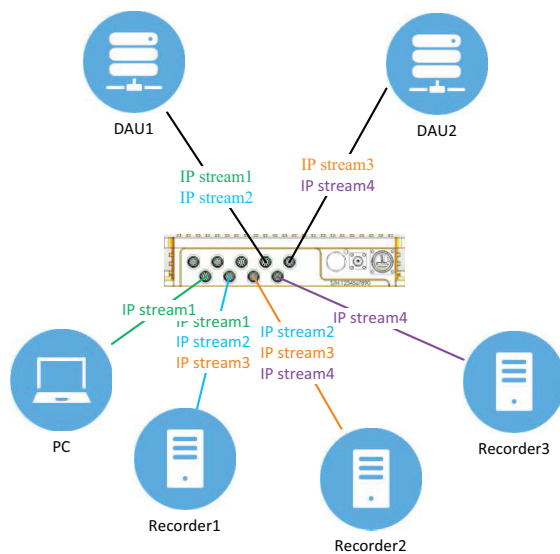


Fig. 5: multicast packet filtering application

In the example shown in Fig. 5, DAU1 and DAU2 send their data as multicast packets. The IP streams from DAU1 are denoted IP stream1 and IP stream2, and the IP streams from DAU2 are denoted IP stream3 and IP stream4. The recorders and the PC will receive all the four streams and each of the two DAUs will receive two streams unless IP multicast filtering is used on the switch.

Four IP multicast filters must be set on the switch in order to ensure that IP stream1 is only received on the PC and Recorder 2, IP stream2 is only received on Recorder1 and 2, IP stream3 is received on recorder 1 and 2, IP stream4 is received on recorder2 and 3 and no multicast stream is received on DAU1 or 2. The IP multicast filters are set on the switch by mapping the four IP multicast addresses to the corresponding four MAC multicast addresses that these IP addresses represent. These four MAC multicast addresses are then written to the switch MAC table with relevant egress port(s) also defined.

Setup of IP multicast filters can be done in two ways:

- static IP multicast filtering
- dynamic IP multicast filtering

### Static IP Multicast Filtering

IP multicast addresses can be setup via Hypertext Transfer Protocol (HTTP), Simple Network Management Protocol (SNMP) or Command Line Interface (CLI).

The user specifies the IP multicast address and the egress port(s), where IP packets with the given IP multicast address shall be forwarded to for each of the static IP multicast filters that he wants to configure.

An IP multicast packet with a destination address not defined will either be dropped at the switch or handled in the same way as a broadcast packet. That means forwarded on all egress ports (except the receiving port) belonging to the same VLAN as the receiving port.

The multicast consumers must be connected to a switch port configured with correct multicast filter settings.

### Dynamic IP Multicast Filtering

The Internet Group Management Protocol (IGMP) standard, see [5], describes a method where multicast consumers can subscribe to one or more IP multicast groups. The multicast consumers makes an IGMP join for each of the IP multicast groups (addresses) that the end node wants to receive data from. These IGMP joins are sent to the querier (i.e. IGMP server) in the network. The IGMP querier is normally a router, but a switch can also act as the IGMP querier if no router is present in the network. The switches can snoop these IGMP control packets, before the IGMP packets are forwarded. Using this information, switches are able to dynamically learn where the multicast consumers are located in the network and then automatically configure multicast filters accordingly so that the multicast traffic only is forwarded to subscribing ports. Non-snooping aware switches simply forward all multicast traffic; including the IGMP control packets, as if it was regular broadcast traffic with no multicast filters as a result.

A multicast consumer must send IGMP joins, while a multicast producer is not supposed to do so unless the multicast producer also wants to receive multicast data sent to a given multicast group.

FTI systems are based on using a high number of multicast groups for data acquisition. The reasons for this are as follows:

1. A multicast group (IP multicast address) is often allocated for each type of data source.
2. Data producers transmit their multicast data to the network. This means that the multicast data is sent even if no data consumer has explicitly requested the data. Data for a given multicast group will be sent to the switch to which it is connected. If IGMP snooping is enabled and no multicast consumer in the network has subscribed to the data from this multicast producer, then the multicast packet for this group will be dropped. However, if one or more multicast consumers exists in the network, then these nodes will notify the switches by sending out IGMP joins, and multicast packets will be forwarded to these multicast consumers.

The user can connect a multicast consumer to any switch port in the network. The switches will automatically detect where the multicast consumer is located and filters will be setup accordingly. The switches will also remove no longer active filters, when a multicast consumer is removed from the network or the multicast application on the multicast consumer is terminated.

This multicast concept may represent a high IGMP control packet load on the switches if the end nodes and the switches are based on IGMP version 1 (see [3]) or version 2 (see [4]) since an IGMP Membership Report for these two versions of the IGMP standard, is based on sending individual IGMP joins for each IP multicast group.

This can be a concern because the IGMP control packet load can represent a performance issue on the IGMP snooping switches if the number of IP multicast groups is high. IGMP version 3 (see [5]) solves this problem since IGMP bulk joins are defined in this version of the IGMP standard. Up to 183 multicast groups can be included in single IGMP version 3 Membership Report (IGMP join). This drastically reduces the IGMP control packet load in the network. If IGMPv1 or IGMPv2 is used, then up to 500 IGMP Membership Report packets can be generated from a multicast consumer each time this end node receive an IGMP Query packet from the network, while IGMPv3 will reduce this high number to only three IGMP Membership Reports since each IGMPv3 Membership Reports can hold up to 183 IP multicast groups.

To avoid overloading network switches, all multicast end nodes subscribing to large numbers of multicast groups should support IGMP version 3. Furthermore all IGMP snooping switches should also support IGMP version 3 as well.

### UDP packet filtering

In a FTI system the Ethernet data is to a large extent based on UDP, see [6].

UDP packet filtering has been introduced in FTI since the number of unique IP multicast addresses used in some case are much less than the number of data streams originating from the data acquisition units. Some FTI systems are even based on broadcast. UDP packet filtering on the switches can be used in order to solve this problem.




Fig. 6: UDP packet filtering application

The setup shown in Fig. 6, is based on a DAU that sends four multicast streams with the same multicast IP address as the destination address. The streams are only differentiated by the UDP source or destination port number. The UDP port numbers can be used as filtering criteria on the switch.

Packet filtering based on UDP source/destination port numbers is not a standard Ethernet switch filtering technique, but this technique has become widely used in the FTI community. When using UDP packet filtering, one must be compliant with the rules defined by the Internet Assigned Numbers Authority (IANA), which is responsible for maintaining the official assignments of UDP port numbers for specific use. Care should be taken to avoid that the switch does not block UDP protocols such as SNMP or Trivial File Transfer Protocol (TFTP).

## Conclusion

The increased number of bandwidth demanding data acquisition units in complex FTI networks requires smart packet filtering techniques implemented in the Ethernet switches in order to handle network bottlenecks and ensure that Ethernet end nodes are not overwhelmed with data not meant for the given node. VLAN (layer 2), Multicast IP filtering (layer 3) based on either static settings or dynamic setup based on IGMP snooping (recommended) or UDP port number filtering (layer 4) are relevant packet filtering techniques for the FTI community.

## Reference

[1] IEEE 802.1: *802.1Q* - Virtual LANs

<http://www.ieee802.org/1/pages/802.1Q.html>

[2] IGMP snooping, RFC 4541,

<http://www.ietf.org/rfc/rfc4541.txt>

[3] IGMP version 1, RFC 1112,

<http://www.ietf.org/rfc/rfc1112.txt>.

[4] IGMP version 2, RFC 2236,

<http://www.ietf.org/rfc/rfc2236.txt>.

[5] IGMP version 3, RFC 3376,

<http://www.ietf.org/rfc/rfc3376.txt>.

[6] User Datagram Protocol RFC 768

<https://tools.ietf.org/html/rfc768>

[7] IANA Internet Assigned Numbers Authority

<https://www.iana.org/>