What Ethernet Switch should I use with GigE Vision Cameras in an embedded system?

Introduction

Embedding high-resolution GigE Vision cameras in drones for applications like pipeline inspection, surveying, and mapping provides valuable visual data, but requires an Ethernet switch that can reliably handle the camera bandwidth and data rates. This paper provides guidance on key criteria to consider when selecting an appropriate Ethernet switch.

Overview of GigE Vision and Key Requirements

GigE Vision is a camera interface standard that uses standard Gigabit Ethernet to deliver high data rates up to (and lately, exceeding) 1000 Mbps per port, making it well-suited for high-definition video streams. However there is often some confusion as to what kind of ethernet switch is necessary to work with GigE Vision. The good news is that GigE Vision is *technically* compatible with any standard gigabit ethernet switch. However, there are some things that need to be considered when actually choosing a switch for this purpose. Transferring image and video data from multiple GigE Vision cameras requires an Ethernet switch with:

1. High bandwidth
2. Enough ports for all cameras
3. Non-blocking switch fabric
4. Sufficiently high packet buffer memory
5. Larger MTU size support (Jumbo frames)
6. Low latency for real-time processing
7. Management features like PTP and Multicast
8. PoE Capability
9. Environmental durability
10. Compact size and low weight

Ethernet Switch Selection Criteria

When selecting an Ethernet switch for deploying GigE Vision cameras on an embedded drone system, here are 10 key factors to consider:

1. High bandwidth

The Ethernet switch needs sufficient bandwidth to handle gigabit ethernet traffic. This is to say, GigE Vision devices required gigabit ethernet (1000BASE-T) ports. While this may seem an obvious point it's worth mentioning that while many modern devices have a gigabit ethernet port, very few of them will actually utilize the full 1Gbps bandwidth available. This is not true of GigE Vision, where a camera can be expected to output a full 1Gbps. With this in mind, it should be obvious that an ethernet switch used with GigE Vision should at least be compatible with 1000BASE-T.

2. Enough ports for all cameras

This is a simple aspect; if you need lots of cameras, you need lots of ports. There is typically a tradeoff between size, weight, power and functionality here, especially on mobile platforms, and typically it is preferable to have as minimal number of cameras as possible. That being said, a typical imaging payload can easily have 3-4 different cameras integrated, not just to achieve different optical FOVs but also to achieve different spectral FOVs.

As such, the ethernet switch must have at least as many ports as cameras, plus one for an upstream connection. Ideally the upstream connection is also a higher bandwidth port such that data can be aggregated from all cameras and transmitted upstream. A common way to achieve this is using a gigabit ethernet switch with a single 10Gbps (10GBASE-T) port for the uplink.

3. Non-blocking switch fabric

With multiple ports operating at bandwidths exceeding 1Gbps, it is crucial that the actual fabric of the ethernet switch is capable of supporting the aggregate data rate of all ports combined. As an example, an ethernet switch consisting of 4 1000BASE-T ports and 1 10GBASE-T port should be able to support at least 14 Gbps on all ports simultaneously to be considered truly non-blocking.

Thankfully, nearly all ethernet switches on the market are non-blocking, as it doesn't make much sense for an ethernet switch to be any other way.

4. Sufficiently high packet buffer memory

Ethernet switches contain internal memory elements (buffers) that are used to store packets temporarily. The reason for a packet buffer is not obvious, especially given the fact that most ethernet switches are non-blocking; if a switch can simultaneously receive and forward data from all ports at their maximum speed, why should any internal memory storage be necessary?

The problem occurs when multiple devices on the switch attempt to send their traffic to a single device. Imagine three devices attempting to send three individual 1Gbps streams in a burst to another device on a 1Gbps on the switch. The receiving device cannot receive 3 Gbps on a 1Gbps connection, and thus packets need to be stored in a buffer temporarily. The size of this buffer determines the maximum duration of this burst, and is a key consideration in selecting an ethernet switch.

As an example, an ethernet switch with a packet buffer of 1.75Mbits could sustain such a burst for around 875us (approximately 72 standard 1500 byte MTU frames stored), whereas an ethernet switch with a packet buffer of 0.096Mbits could only sustain such a burst for 48us (approximately 4 standard MTU frames stored) (see calculations below).

Buffer: 1.75Mbits

3 x 1000 Mbits = 3000 Mbits / s

MBuffer = 1.75Mbits

If we assume that 1000Mbits can be simultaneously streamed, then 2000Mbits is the incoming data rate.

Fdatarate = 2000Mbits

$$T_{\text{burst}} = \frac{1.75 \text{ Mbits}}{2000 \text{ Mbits/s}} = 875 \, \mu\text{s}$$

With a packet buffer of 0.096Mbits…

$$T_{\text{Burst}} = \frac{0.096 \text{ Mbits}}{2000 \text{ Mbits/s}} = 48\ \mu\text{s}$$

Thus, buffer size directly affects the length of sustained “many-to-one” bursts that the ethernet switch can handle and is an important consideration when building a GigE Vision based network.

5. Larger MTU size support

An Ethernet MTU (sometimes referred to as a packet or frame) is the largest discrete group of bytes that is sent in a single transmission. Typical ethernet limits this to 1500 bytes in total, with 40 bytes of header and 1460 bytes of payload.

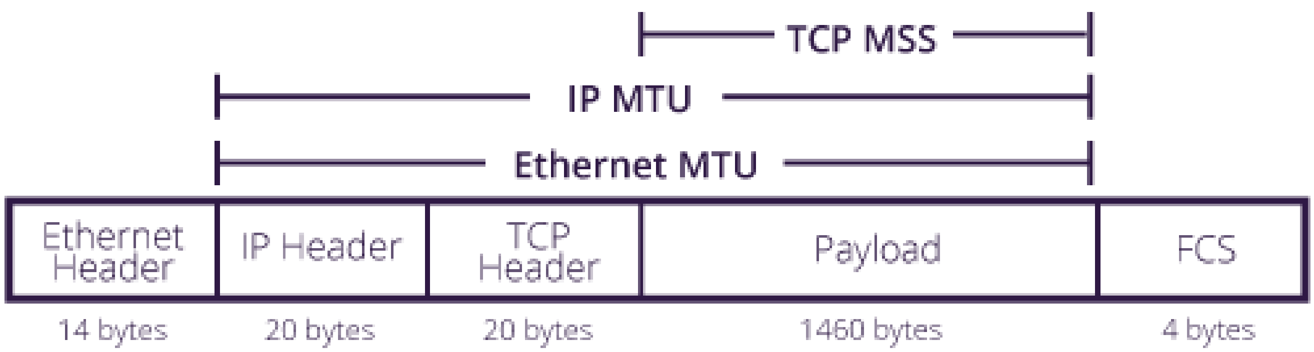

A small MTU is beneficial because it reduces the chance of an error causing frame corruption ; essentially, less time on the wire means less likelihood of bit errors). However small MTUs mean that less of the bits actually sent are payload, meaning reduced efficiency and reduced overall data rate. For this reason, many GigE Vision cameras are designed to support an MTU larger than 1500 bytes, the most common being 9000 bytes (called Jumbo Frames).

Not all ethernet switches support an MTU larger than 1500 bytes, and thus it is crucial to check this before purchasing. Some ethernet switches state that they support jumbo frames but require some configuration first, so it’s very important to figure out exactly how this configuration works before making a purchase decision.

Note, it is usually possible to run GigE Vision cameras using standard 1500 byte MTUs, so even if a switch does not support jumbo frames, it may still work for your application. However you won’t be getting the most out of the GigE Vision camera in this case.

6. Low latency for real-time processing

In some vision applications, latency is a huge determinant of system performance. Ethernet switches introduce an inherent latency which can reduce overall system performance.

There are two main switch architectures that affect latency; store-and-forward, and cut-through. In both cases, the role of the switch is to read the destination MAC address of an ethernet frame, and send that frame out onto the correct port (based on the switch's internal MAC address table). In store-and-forward mode, the switch receives the packet in full, stores it, reads the MAC address, then forwards the packet onwards. In cut-through mode, the switch need only receive the frame up to the destination MAC address part (the first 6 bytes of the ethernet header), at which point it can start transmitting out the packet immediately.

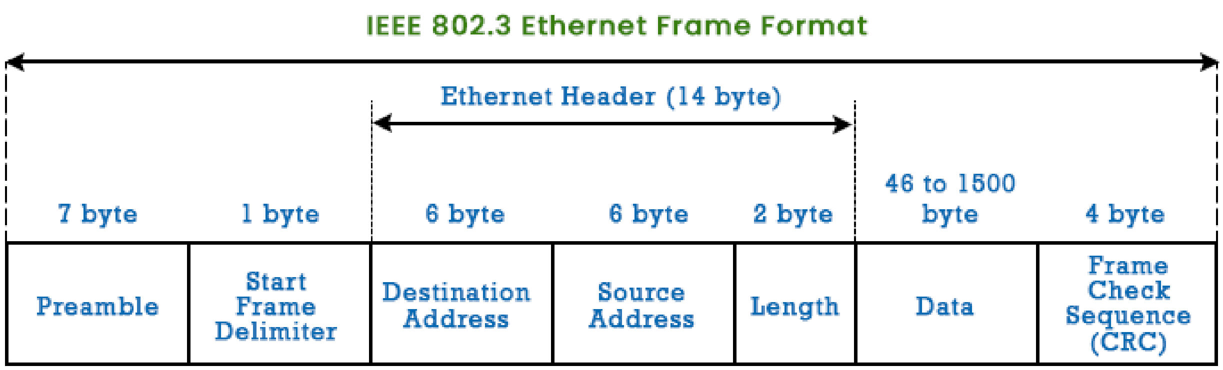

As should be evident, cut-through mode means the packet spends less time in the switch, and thus reduces latency. In terms of numbers, one may expect a 9000 byte frame to have a latency of at least 72 us under store-and-forward, and 48 ns under cut-through (calculations below).

$$T_{\text{read}} = \frac{\text{Bits}}{F_{\text{datarate}}}$$

Fdatarate = 1000Mbits/s

Bits = 9000 bytes = 0.072Mbits

$$T_{\mathrm{read}} = \frac{0.072\ \mathrm{Mbits}}{1000\ \mathrm{Mbits/s}} = 72\ \mu\mathrm{s}$$

Fdatarate = 1000Mbits/s

Bits = 6 bytes = 0.000048Mbits

$$T_{\mathrm{read}} = \frac{0.000048\ \mathrm{Mbits}}{1000\ \mathrm{Mbits/s}} = 48\ \mathrm{ns}$$

These calculations are rough minimums, and should be at least doubled to account for extra PMA level latencies (SerDes, Scrambling, etc).

This all being said, the latency of an ethernet switch, even in the worst case, typically isn't going to exceed 200us under store-and-forward, which in most real world applications is "good enough". It's worth noting that latency of a vision system is usually far more constrained by the embedded computer actually doing the processing, rather than the switch and network infrastructure.

7. Management features like PTP and Multicast

The beauty of GigE Vision is that it builds upon the existing technologies of gigabit ethernet, meaning all the other features of ethernet can be used alongside Gige Vision to achieve efficient and sophisticated networks. Thus, when it comes to selecting an ethernet switch, it's important to consider which of these features the switch supports.

There are three common ethernet technologies that are typically used alongside GigE Vision.

1. IGMP Multicast - Allows for the efficient use of network bandwidth by allowing only interested subscribers to receive data
2. VLAN - Allows networks to be logically segregated for security and network reliability (for example, segregating image data and critical control data in a single switch)
3. IEEE1588 PTP - Allows precise timing to be achieved using ethernet packets

The three features above are by no means exhaustive, but are typically incredibly useful in efficiently using network infrastructure (and efficiency is a key factor on mobile platforms where you can't simply throw another 1U rack switch at the problem to solve it).

PTP is also a hugely advantageous technology especially when it comes to imaging applications that are trigger based, where accurate timing down to a 1 microsecond timebase is necessary.

The vast majority of embedded switches do not support much in the way of management, and those that do are typically poorly documented and run unreliable, unmaintained software. Therefore many applications have to settle for larger enterprise networking solutions to build such networks.

8. PoE Capability

On a mobile platform, less wires = less weight, and many GigE Vision cameras will support Power over Ethernet (PoE), which allows the same wires used for ethernet data to also transmit power. Therefore it's advantageous to use a switch that can support PoE. Pay close attention to the type of PoE required by your camera, and the type supported by the switch. The latest PoE standard (IEEE 802.3bt) supports up to 90 W per board; not all switches will support this.

9. Environmental durability

When considering operation on a mobile platform (which can often take the form of a drone, vehicle or naval robot), environmental durability and resistance to heat, cold, vibration and shock are key considerations. The typical enterprise network switch is designed for operation in a data center with controlled temperature, airflow and minimal movement. Standard commercial home network switches fare a little better but are still ill suited to operation in extreme environments.

Therefore when selecting a switch for use with GigE Vision, it's important to consider environmental durability. Look for devices that use vibrational tolerant connectors, use automotive rated semiconductors, and systems that have been qualified and tested under environmental testing schemes like MIL-STD-810G. Additionally empirical and calculated MTBF figures from the manufacturer can be highly useful in further ascertaining the suitability of an ethernet switch for an embedded vision application.

10. Compact size and low weight

During integration, size and weight are key considerations for mobile platforms, and especially so for vision systems on UAVs, where internal volumes are limited and every gram matters. Thus it's crucial to use a switch that is optimized for low size and weight.

Conclusion

Ultimately, only thorough real-world testing is able to truly confirm that an ethernet switch is suitable for a GigE Vision application. This is especially true when considering that the ethernet switch needs to be rugged and small, while still being reliable and potentially capable of sophisticated network topologies.

Choosing the right Ethernet switch is essential for building robust drones with GigE Vision cameras. This whitepaper has provided guidance on key selection criteria - leverage this information to evaluate switches and ensure they fulfill bandwidth, durability, footprint, and integration requirements.

Disclaimer

Technical specifications may be subject to change without notice. All trademarks are acknowledged as property of their respective owners. BotBlox cannot be held responsible for any damage in equipment or subsequent loss of data or whatsoever in consequence of this document.