Switches for Networked FTI

Nikki Cranley, Ph.D, Diarmuid Corry, M.Sc.
ACRA CONTROL INC., Maryland, USA

ABSTRACT

Ethernet technology offers numerous benefits for networked Flight Test Instrumentation (FTI) systems such as increased data rates, flexibility, scalability and most importantly interoperability owing to the inherent interface, protocol and technological standardization. In a networked FTI system, the switch is a key component that allows data to be routed between Data Acquisition Units (DAU's), networked recorders, data processing and analysis stations. This paper provides an introduction to network switching concepts with a focus on its operation in a networked FTI system. The features of Commercial Off-The-Shelf (COTS) and FTI switches are compared demonstrating the benefits of FTI switches in terms of reliability, routing, throughput, latency, and start-up delays.

1. INTRODUCTION

In recent years there has been a shift from proprietary and closed solutions for Flight Test Instrumentation (FTI) networks towards more open standards-based systems using Ethernet technology. The trend towards Ethernet is further driven by the CTEIP iNET initiative that is pushing the adoption of Ethernet technology for the future of FTI. Using open standard Ethernet technologies offers the FTI community greater flexibility in system design and more choice for multi-vendor interoperable systems. Some successful networked FTI systems have been published in [1,2, 3].

The adoption of Ethernet technology brings about an important paradigm shift from traditional FTI affecting all aspects of the DAU design, operation and the system topology. The more significant aspects of this paradigm shift include:

- **Master/Slave to DAU independence** - In traditional FTI systems comprising two or more DAU's, one DAU is typically designated to be the master to the remaining slave DAU's. The master DAU performs key functions: It acts as a gateway or conduit through which slave DAUs may be accessed; is responsible for synchronizing the slaves ensuring sampling coherency; gathers and places data from itself and from the slave DAUs for transmission as PCM. In contrast in networked FTI systems there is no need for a specific master. Each networked DAU operates independently of all other nodes in the system. The networked DAU is responsible for its own data acquisition, data transmission, and synchronization to a common time source.

- **Determinism** - PCM systems are deterministic such that it was possible to determine the propagation delay for a sampled parameter to reach its destination. This determinism allowed modern synchronized systems to guarantee that all data in a PCM major frame was sampled in the same epoch (coherent), so only the time stamp of the frame was required in order to calculate the sampling instant of each parameter without reference to the frame layout. In contrast, Ethernet IP networks are non-

deterministic. Ethernet networks lack deterministic propagation delays, which means that although a parameters' sampling instant may be time stamped, the actual arrival time of this data at its destination is known only with probabilistic accuracy. Such uncertainty creates added complexity when quantifying real-time latency budgets, or when placing networked data into a deterministic legacy PCM stream.

- **Start up delay** – Devices in a traditional FTI system are capable of becoming "live" and transmitting data within a short interval (< 1sec). Brown outs or power failures should have a minimal effect on data acquisition. And in contrast in networked FTI systems, due to the increased complexity and added functionality of the network nodes, this start up delay may be increased by a factor of 10 or more. Upon power-up, Ethernet technology allows network nodes to discover, perform dynamic self-configuration, and learning of the nodes configuration and topology.

In an FTI network of distributed peer DAU's, the switch is a key component that allows data to be transmitted to and from different nodes in the network. Switches comprise a number of ports to which other networked devices may be inter-connected such as DAUs and switches. In general, connections in the switch utilize point-to-point full-duplex Ethernet links. By far, the most important task of the switch is to reliably and quickly forward packets to their destination. The packets received by the switch are stored in a buffer until they reach the head of the queue. Once at the head of the queue, the switch examines the packets' destination and through a lookup mechanism determines how to forward the packet to its intended destination. This process is known as Store-and-Forward since the packets are stored in a queue until they are switched or forwarded. It is this store-and-forward mechanism that is the primary contributor to the best-effort nature of Ethernet in terms of:

- Queuing delay and switching delays that introduce latency and jitter
- Packet loss and buffer overflow
- Misordered packets
- Compromised effectiveness of IEEE 1588 Precision Time Protocol (PTP) if it is not supported in the switch
- Learning the network topology and power-up delays

FTI switches are designed to minimize these best-effort side effects of Ethernet technology.

Switches vary in complexity: unmanaged switches have no configurable options while managed switches have a configuration interface and additional management features. Although the features available on a switch is largely vendor-specific, typical managed switch options include:

- Provide a bridging mechanism to higher speed Gigabit Ethernet links or to other networking technologies, such as Wireless LAN.
- Prioritization of packets based on IP address, IP priority flags using the Differentiated Service Code-Point value (DSCP) or MAC layer prioritization using VLAN 802.1q tags;

- SNMP Agents to report network-related statistics e.g. packets received per second, packet loss etc.;
- Sniff and Mirror ports allowing the packets on a given interface to be copied to another interface for analysis.

The remainder of this paper is structured as follows: Section 2 describes the operation of a generic store and forward COTS switch, Section 3 outlines the requirements of a switch for FTI in terms of physical characteristics, time synchronization support, and routing. Section 4 provides a summarized comparison of generic COTS switches and FTI switches.

2. FUNDAMENTALS OF LAYER 2 STORE-AND-FORWARD SWITCHING

This section provides an overview of generic layer 2 switching concepts. First, it is necessary to consider the operating environment for which a generic Layer 2 Ethernet switch was designed. Generic switches are designed for environments where:

- Devices may be plugged into the network at random
- Devices may dynamically acquire IP addresses from a DHCP server
- Devices require data to be routed to a variety destinations that may vary over time
- Devices may have on-off transmissions to a variety of destination devices.

For these reasons a generic switch must be capable of learning the routing information dynamically over time. Although, from an FTI perspective the dynamic learning feature of generic switches encumbers deterministic behavior, consumes additional processing/switching resources, and typically have a longer start-up delay.

There are two classes of Layer 2 switches: Store-and-Forward and Cut-Through switching. The reader should be cognizant that although the internal operations in the switch are not standardized with respect to how the packets are switched, the following provides a high-level overview of the switching implementation in a generic COTS switch.

2.1. STORE-AND-FORWARD SWITCHES

The packets received by the switch are stored in an input buffer until they reach the head of the queue. Once at the head of the queue, the switch core buffers the entire Ethernet frame and performs error checking on the Ethernet frame by comparing the Ethernet frames FCS against the CRC calculated by the switch core. If the Ethernet frames own FCS differs from the calculated CRC the frame is considered to contain physical or data-link errors and is dropped. In this way, the corrupt Ethernet frame is prevented from propagating through the rest of the network.

If the Ethernet frame is determined to be valid, the switch core examines the packets' destination and through a lookup mechanism in the MAC address table to determine how to forward the packet to its intended destination. If there is no entry for a given Destination MAC address, the switch does not know where to forward the packet. In this case, the Ethernet frame is forwarded out all ports on the switch, or flooded. As Ethernet frames are passed through the switch, the switch core updates the MAC forwarding table noting the

Source MAC address and the interface on which it arrived. By doing this, the switch core is able to dynamically populate the MAC forwarding table, however since MAC tables have a finite memory size, entries age out to ensure that the table is up to date.

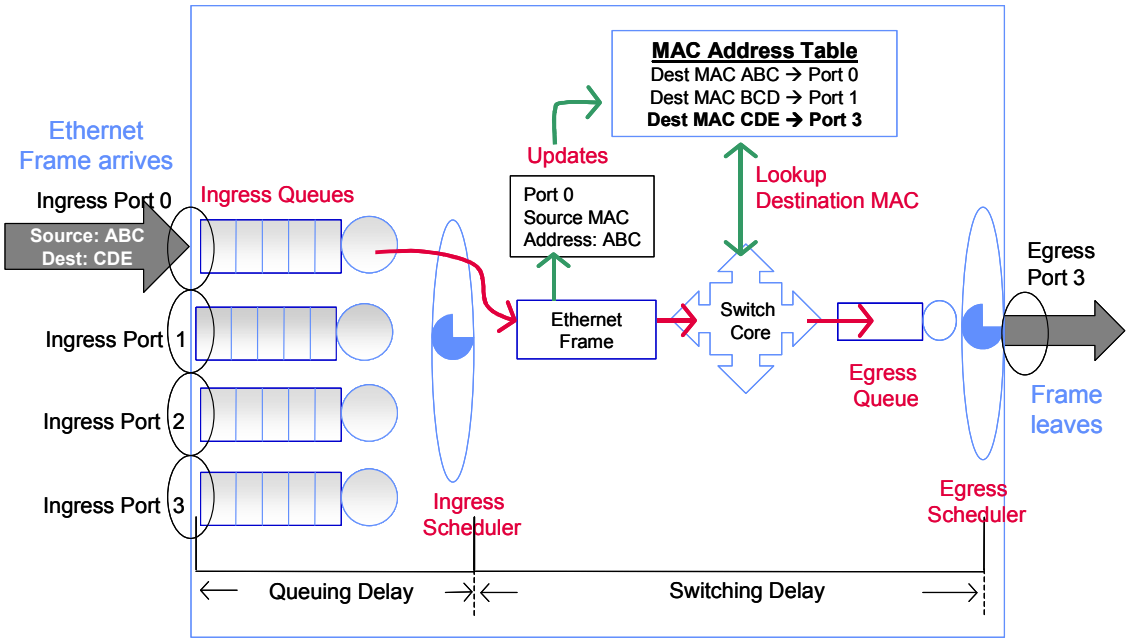

Figure 1: Generic COTS Switch

For example, consider an Ethernet frame with a source MAC address of ABC and Destination MAC address of CDE arrives on Port 0 of a generic Ethernet switch as shown in Figure 1. The Ethernet frame is buffered in the ingress queue until it reaches the head of the queue. Once at the head of the queue, the ingress scheduler passes the Ethernet frame to the switch core where error-checking operations are performed. The switch core takes note of the port on which this Ethernet frame arrived and its source address, so that should any Ethernet frame need to be forwarded to this destination i.e. ABC, the switch knows that the destination can be reached through Port 0. Similarly, to determine the forward path of the Ethernet frame, the switch core uses the Destination MAC address, CDE, to lookup the egress port of Ethernet frame i.e. Port 3.

2.2. Cut-Through Switches

In contrast to store-and-forward switches, cut-through switches do not need to buffer the entire Ethernet frame before making the forwarding decision. Rather, a cut-through switch examines only the Destination MAC address before the switch core is capable of determining the outgoing egress interface. By examining only the first 6Bytes of the Ethernet frame, cut-through switches are faster than store-and-forward switches. This can result in considerable savings where there may be jumbograms that need to be switched in the system. This is due to the fact that it takes approximately the same amount of time to switch an Ethernet frame regardless of size by virtue of the fact that only the first 6Bytes are examined before the forwarding decision can be made. For example, consider a 9000Byte Ethernet frame, a cut-through switch can determine the forward path for this frame having read only the 6Bytes of the Destination MAC address. In contrast, a Store-and-Forward switch would need to buffer

9000Bytes of the Ethernet frame, perform error checking to verify the FCS of the Ethernet frame with the switches own calculation of the CRC, finally the switch core examines the Destination MAC address to determine the outgoing egress interface. This can result in a time saving of a few microseconds using cut-through switching as opposed to several milliseconds using a store-and-forward approach. However, although there is a significantly reduced switch since there is no error checking of the Ethernet frame, cut-through switches do not drop invalid Ethernet frames resulting in the propagation of the corrupt Ethernet frame through the rest of the network, thereby violating the IEEE 802.1D bridging specification. However, typically the receiving device discards corrupt Ethernet frames.

2.3. SWITCHING QUALITY OF SERVICE ISSUES

Generic layer 2 switching introduces delay and jitter to the packets as they traverse the switch. The switching delay comprises two contributing components, the queuing delay and the actual switching delay. The queuing delay is the waiting time of the Ethernet frame to reach the head of the queue before the switch core takes over. The queuing delay is a function of:

- The occupancy of the input queue both in terms of number of packets and packet size
- The queue length
- The input queue scheduling algorithm
- The number of input queues - The number of queues rather than the number of ports is considered since QoS mechanisms may allow pre-classification and differentiated services of the input traffic prior to switching.
- The switching delay occurs during the error detection and Destination MAC lookup mechanism.

These variables impact on the accuracy of distributed time synchronization, reliability, and the ability to plan the network for deterministic behavior.

3. SWITCHES FOR FTI

Although Ethernet technology offers the ability to use off-the-shelf networking devices, FTI networked systems have specific requirements that are not typically met using standard COTS switches. Unlike a generic network, FTI networks have strict Quality of Service targets and design goals for acceptance in terms of reliability, determinism, and latency, however the unique nature of FTI networks can be exploited to optimize switching functions.

3.1. REQUIREMENTS OF FTI SWITCHES

The FTI network topology is static; nodes are not dynamically added/connected to the network. In contrast, generic COTS switches are used in environments where almost any device can be plugged in anywhere in the network. Typically FTI networks employ a cascading switch tree-type topology which is adopted with full-duplex DAU-switch and inter-switch connections where the number and types of networked nodes in the system needs are known (e.g. DAU's, network recorders, 1588 grandmaster, on-board data processing stations and PCM RF link requirements). As a consequence of the static topology of the FTI

network, the aggregate data rates generated by each DAU impact on the choice of link speeds, switch-interconnection, and physical media, i.e. fiber or copper Fast Ethernet or Gigabit Ethernet.

Unlike the operating environment of a generic COTS switch, FTI networks are carefully designed and planned. For each DAU the number of data flows and the intended destinations for each flow are known. Furthermore, not only is the per-flow mean bitrate known, but also the characteristics of the data-flow in terms of its packetization and packet rate. In this way, FTI networks have little, if any, dynamic traffic, variable bit rate data or adaptive protocols on the network during data acquisition. Moreover, during data acquisition there are few dynamic multicast joins (if IGMP is enabled). Assuming stable and constant traffic conditions, transmitting data in a deterministic manner results in quasi-deterministic output data from the switch. Coupling knowledge of the FTI switching behavior and knowledge of the peak aggregate data rates, analysis can be used to determine the upper-latency bounds to be calculated and the network can be designed for reliability and congestion avoidance by over-provisioning to guarantee the aggregate peak data rate on each link never exceeds the capacity of that link.

An important feature of FTI switches is that they should be live-at-power up within several seconds and recover quickly and faultlessly from electrical brownouts. Generic COTS switches require a certain degree of "dynamic learning" when they are powered on. For example, switches with dynamic routing must keep a record or MAC table of which source and destination IP addresses are connected to the respective interfaces. From the MAC table the switch learns how to route packets from a given source to destination. Not only does this learning take time but it must also be performed continuously so that the MAC table is fresh since sources and destinations may change over time. Since the FTI network topology is static, the desired routing for the data is known. FTI switches typically support hard-wired routing tables or for the routing table to pre-programmed into the switch. Due to the "simplified" hard-wired routing, the switches behavior in terms of queuing and switching algorithms can be known, thereby allowing for mathematical analysis. Moreover, should brown-outages occur, these routing tables are retained and maintained on power-up thereby allowing for faster routing, faster startup, and faster recovery from brownouts than those with dynamic "learning" features.

3.2. 1588 SUPPORT IN SWITCHES

IEEE 1588 PTP Time support essential in a distributed networked FTI to ensure coherency and simultaneous sampling. The unusual aspect of FTI networks is that the traffic load in the network is heavily asymmetric particularly at medium to heavy data rates. This asymmetry arises from the fact that the DAU's acquire and transmit data at a much faster rate than they receive data. This affects the efficacy of the network to carry PTP messages between the grandmaster and the DAU's since one of the underlying assumptions of 1588 PTPv1 is end-to-end network delay symmetry. To overcome this asymmetry and improve synchronization accuracy the switch should support 1588 compensation mechanisms. The two compensation mechanisms include 1588 Boundary Clock and 1588 Transparency.

In Boundary Clock mode the switch itself is synchronized in slave mode to the attached master clock, shown in Figure 2. Once synchronized the switch then acts as the master clock to all the attached slaves. In this way, delay asymmetries and internal delays are compensated. The disadvantage of Boundary Clock mode is that synchronization is cascaded

in that synchronization must take place at each switch hop between the Grandmaster/Master and the Slave resulting in an accumulation of non-linear time offsets thereby degrading the synchronization accuracy. This is particularly problematic in highly cascaded network topologies.

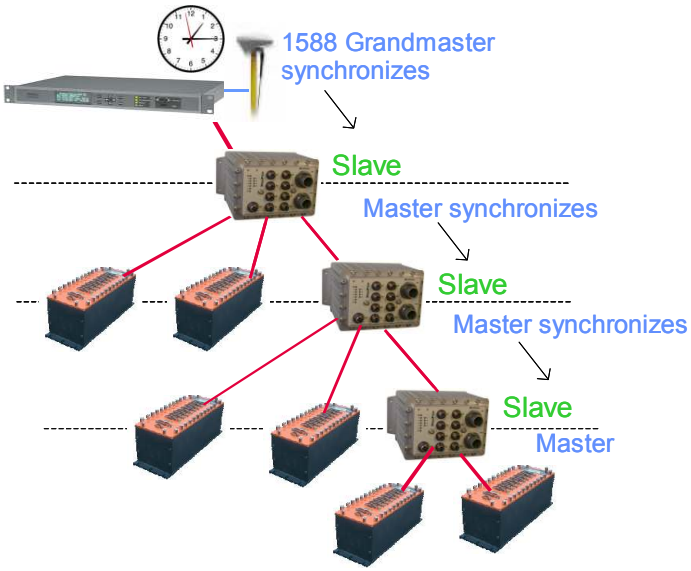

Figure 2: Boundary Clock Operation

Transparency is not part of the IEEE 1588 PTPv1 standard and is therefore implemented in a vendor-specific way for PTPv1. However, there are significant commonalities in the various transparency implementations. Let D be the propagation delay for the 1588 PTP packet between hops and let Q be the time in residence (i.e. the queuing and switching delay) experienced by the packet in the switch. Without transparency, PTP packets experience propagation, queuing, and switching delays with each hop from Grandmaster to Slave, shown in Figure 3. Since Q is unknown, this leads to an error associated with the timestamp embedded in the 1588 PTP packet.

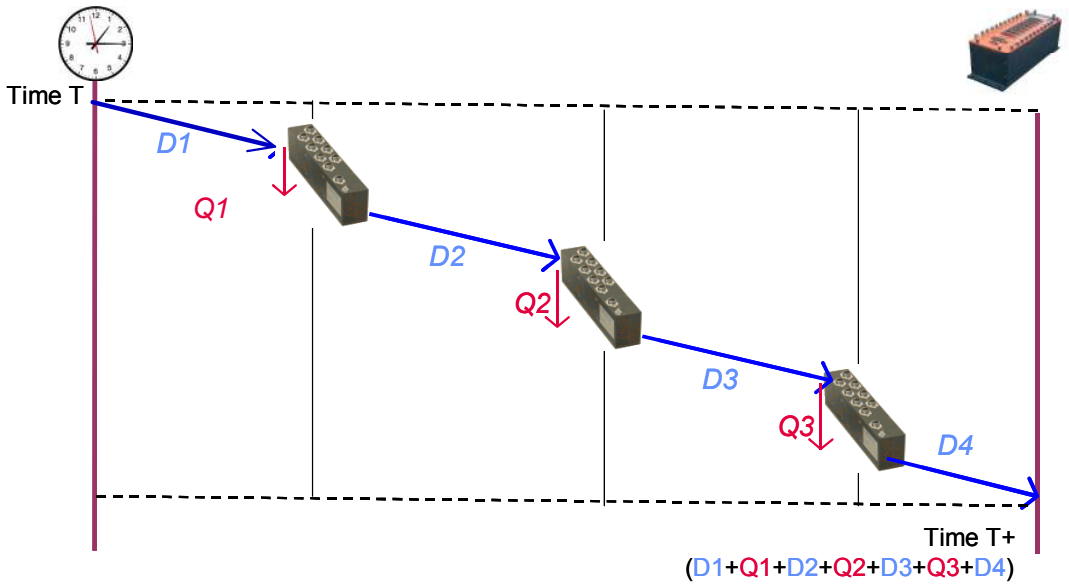

Figure 3: Switching delays without Transparency

In Transparency mode, the switch intercepts the incoming PTP packets containing an embedded timestamp, T. The switch measures the time in residence, Q, for a given PTP packet. Before the PTP packet is forwarded on, the timestamp, T, embedded within the PTP packet is modified to compensate for this residence time in the switch. Effectively this renders the switch "transparent" in the network for PTP traffic, as shown in Figure 4. For this reason, transparency results in more accurate end-to-end synchronization between Grandmaster and DAU than can be achieved using Boundary Clock operations.

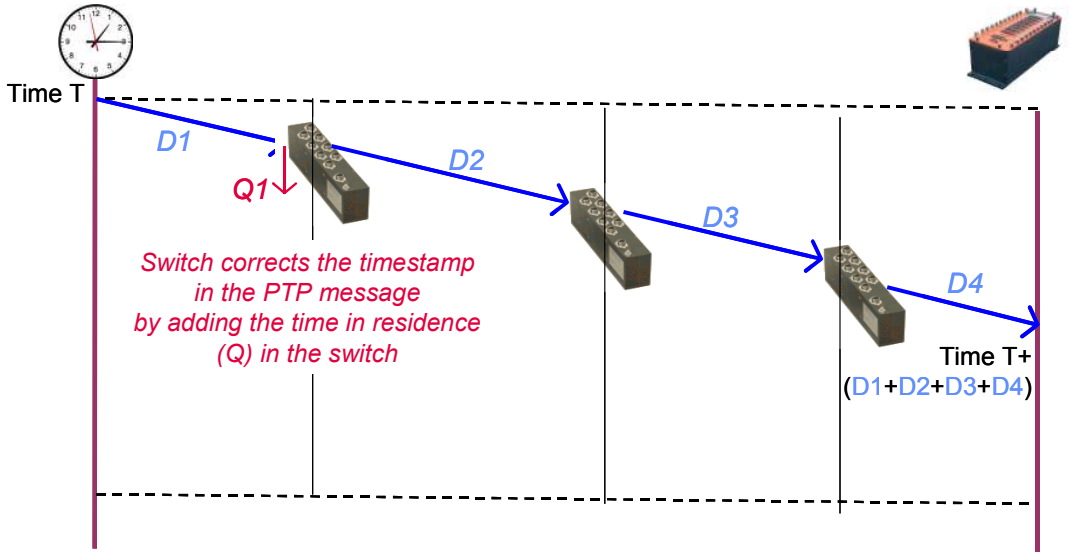

Figure 4: Switching delay compensation with Transparency

4. CONCLUSION

In recent years, there has been considerable momentum in the FTI domain towards the adoption of open standards-based systems using Ethernet technology in recent years. A key-driving factor that facilitated this change was the development of the IEEE 1588 Precision Time Protocol (PTP) for distributed time synchronization. Each element in the network must therefore support for 1588 PTP. A core component in an FTI network is the switch, which routes data from DAU to destination.

This paper provided an overview and introduction to network switching concepts. Although generic COTS switches may operate in an FTI network, they were not designed or optimized, nor meet all the technological requirements of FTI networks. By way of conclusion, the following table provides a summarized comparison of COTS generic store and forward switch and a switch designed for FTI.

Feature	Generic Switch	FTI Switch	Comments
Physical			
Ruggedized 100BaseT Ethernet connectors	N	Y	**Generic:** typically uses RJ45 **FTI:** Mighty Mouse or similar connectors for FTI. Only 4-pins required for 100BaseT Fast Ethernet, 8 pins required for 1000BaesTX Gigabit
Ruggedized enclosure	N	Y	**Generic:** standard off the shelf switches or even rack mounted switches are not always suitable for FTI.
Aircraft power	N	Y	
Industrial Temperature range (-40C/+85C)	N	Y	
Shock and vibration	N	Y	
Ruggedized I/O connector	N	Y	**FTI:** Mighty mouse or similar IO connector with pins allocated for IRIG-time, analog IRIG-time, PPS out, NMEA, CBIT.
SMA Connector for antenna	N	Y	**FTI:** Required as a time source via GPS for the 1588 Grandmaster.
PTP master external 1PPS output	N	Y	**FTI:** 1PPS outputs to verify PTP time synchronisation.
Live at power-up	N	Y	**Generic:** Typically switches have a dynamic learning functionality to allow the MAC tables to be dynamically and continuously populated while the switch is live. **FTI:** Since the FTI Switch has hard wired routing there is no need for dynamic learning of the networked devices in the system.

Mirrored outputs	P	Y	**Generic:** Partial support. Typically called "sniff and mirror" operation. Although typically only a single mirror is possible. **FTI:** Designed for FTI with multiple mirrored ports for data to be recorded, transmitted to an Ethernet-to-PCM converter and also processed on a PC with real-time analysis software.
Core			
PAUSE and Backoff Flow control	Y	NA	**FTI:** This is specifically not suitable for FTI. Any flow control will result in packet loss at the DAU. In a properly designed FTI network, congestion and therefore loss should not occur due to careful planning of the network architecture.
Time			
Internal/External Grandmaster	N	Y	IEEE 1588 2002 PTP v1
Time Source	P	IRIG, GPS	**Generic:** Often switches might have an embedded NTP or SNTP timeserver seeded from a local clock.
IEEE 1588 v1 Transparency	N	Y	**FTI**: Transparency results in better time synchronization in a distributed FTI network.
IEEE 1588 v1 Boundary clock	N	Y	**FTI**: Although transparency results in better time synchronization in a distributed FTI network, some FTI switches implement Boundary Clocks.
Routing and switching			
Fast switching (Fast Ethernet)	N	Y	**FTI**: Due to the hard wired routing there is extremely fast switching of the Ethernet packets (>80Mbps Fast Ethernet) **Generic:** As there is per-packet lookup to reference the MAC table, generic switches are slower than hard-wired routing (<<80Mbps)
Static programmed/hardwired routing table	N	Y	**FTI:** Retained on power down/up. Mitigates the need for dynamically maintained switching tables.
Finite dynamically learned MAC table size for unicast and multicast	Y	NA	**Generic:** The switch would need to continuously update and populate the MAC table resulting in occasional broadcasts of the data when there is no corresponding entry in the MAC table for a given destination. This information is typically not retained on power-up.

DiffServ Quality of Service	P	NA	**FTI:** Not required, as this feature would reduce the switching speed. This would also require a homogenous network whereby the same QoS policy and same switch is used throughout the network infrastructure. QoS is only effective when the network is close to saturation. In a properly designed FTI network saturation and congestion should be avoided as part of the networks design.
Protocols			
IGMP v3 / IGMP v2	Y	NA	**FTI:** Since this has hard-wired routing, this feature is not necessarily required.
SNMP v2c with iNET MIBS and FTI specific MIBS	Y	Y	SNMP v2c RFC1901; SNMP v3 RFC2571. Managed devices collect and store management information and make this information available to the management system using SNMP.
SNTP and NTP Time Server	P	NA	**FTI**: This has 1588 support as neither SNTP or NTP are suitable for FTI.
P: Partially true but not common **Y**: True, supported **N**: Not supported **NA**: Not Appropriate The information is presented based on a review of general switch specifications it is not a comparison against a specific model. Information is subject to change and for specific features.			

It is clear from the Table above that there are many features typically implemented in a generic COTS switch are not required for the FTI domain. Moreover, these generic switch features (such as dynamic learning, dynamic protocol support, adaptive behaviour etc) impede the reliability and performance requirements of networked FTI systems. Furthermore, a generic switch does not necessarily support vital technologies and protocols required by an FTI network, most notably IEEE 1588 PTP. FTI switches have been designed and optimized for FTI networked data acquisition systems implementing standard core switching functions and supplementing FTI specific technological requirements.

REFERENCES

[1] Eccles, Lee, "Network Based data Acquisition", Proceedings of European Telemetry Conference, Munich, Germany, 2008

[2] Revaux, Nathalie and Abadie, Frédéric, "A380 Flight Test Architecture: Switched IENA from Pulse Code Modulation to Ethernet LAN", Proceedings of European Test and Telemetry Conference, Toulouse, France, 2003

[3] Doyle, David and Canizares, Ruben, "Review of a successful distributed networked and modular FTI implementation", to be published in Proceedings of Society of Flight Test Engineers Symposium 39, 2008