Will time-sensitive networking for Ethernet revolutionise industrial networks?

How Time-Sensitive Networking (TSN) for Ethernet Delivers Benefits for Industrial Applications

Introduction

Looking at the connectors on PCs and laptops today and comparing them with devices from just a decade ago, it is incredible to see how standardised everything has become. The slimmest, ultra-thin laptops can now basically get away with a few USB-C connectors and perhaps an audio socket for some headphones. Reflecting on past devices, they would have had power sockets, Ethernet ports, COM ports, HDMI and VGA for video, along with slots and docking station connectors. Wireless technology, such as Wi-Fi and Bluetooth, also negate the need for some of those wired connectors but, overall, it is a lot simpler than it has ever been.

On the industrial networking front, things are very much still like an old laptop; a different connector and protocol seemingly tuned to every sensor and actuator use case. However, things are changing, albeit at the speed industrial engineers are accustomed too – slowly. Ethernet has been a great success for computer networks, simplifying connectivity and reducing configuration challenges while attaining exceptional data rates and throughputs. However, latency and functional safety were never a core focus of Ethernet for most users, hence why it has never really established itself in industrial networks.

This is now starting to change with various working groups having formed and tackled the various latency and bandwidth reservation limitations of Ethernet. This has led to some new standards that deliver benefits not only for industrial networks, but for automotive and audio-visual applications too. Here we will be looking at the current industrial networking landscape, reviewing the new Ethernet standards, and examining the benefits they will bring to industrial systems.

Proprietary is/was King!

Up until the mid-1980s, sensors and actuators would have been read and controlled using analogue technology. Systems such as the 4-20mA current loop were capable of powering the end device while also implementing the communication. It is also noted for its high noise immunity and support for long cable lengths, essential in the implementation of industrial plants where sensors and actuators can be mounted over very long distances to machines, vats and tanks. Failures are also quickly detected. For example, should a cable break, the break in the current loop is easily detected thereby highlighting the problem.

One of the challenges with current loops was the need to run a separate cable to each and every sensor and actuator. With costs of processor technology dropping, along with the integration of serial interface hardware into microcontrollers, multi-drop interfaces using digital communication started to be considered. RS-485 was such an option, utilising the same in-chip peripheral as implemented for RS-232. Combined with the Modbus protocol this was capable of addressing up to 254 devices on one data link[1]. Also using the same physical layer is PROFIBUS, with two variations; Decentralised Peripherals (DP) supports sensors and actuators via a centralised controller, while Process Automation (PA) is designed for use in explosion/hazardous areas[2]. Yet another alternative based upon the same physical layer is CC-Link[3].

Coming from the automotive industry, the Controller Area Network (CAN) bus also provides a robust physical networking interface that is suitable for industrial networks. DeviceNet is an application layer protocol that sits on top of the CAN physical layer[4], while CANopen is another alternative[5].

In what could be termed the Generation I of digital industrial networks, these technologies combined available robust physical layers (OSI model layer 1) together with differing data link layers (OSI model layer 2), adding various amounts of definition to the layers above where required (figure 1). The characteristics of all these networking technologies is their ability to provide robust data transfers over long distances (hundreds of metres to kilometres), defined latencies and, in some cases, capabilities such as determinism and fulfilment of safety requirements.

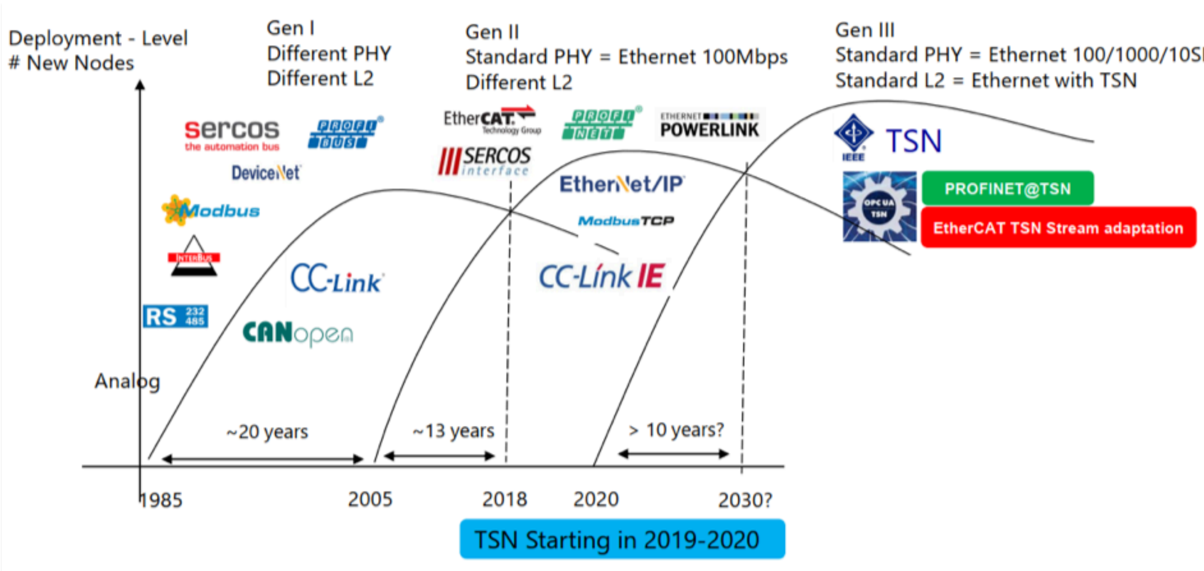

Figure 1 – The deployment of the first three generations of digital industrial networking technology showing gradual convergence toward international standards.

All of those fieldbus technologies listed, plus many others, continue to be in use with millions of nodes deployed worldwide. However, the physical layers used limit data transfer rates to lower than 1Mbit/sec. With the exception of CAN that has seen speed improvements with the introduction of CAN FD, future networking needs needed an alternative technology to handle the larger amounts of data demanded by new industrial systems, especially vision systems. With different physical layers in place, it also meant that the wiring could not be shared across various systems, i.e. CANopen can't transport data over a PROFIBUS network.

Ethernet is an ideal alternative with the physical layer being well established, offering high data rates, and silicon technology being broadly available on the market. Ethernet-based networking technologies started coming to market, establishing Generation II industrial networking technology. These are considered to be the solutions using a common Ethernet physical layer (layer 1) but with differing approaches to the link layer (layer 2) (figure 2). With solutions being introduced around 2005 it required around 13 years for these Ethernet-based technologies to attain parity of node installation with Generation I industrial networks.

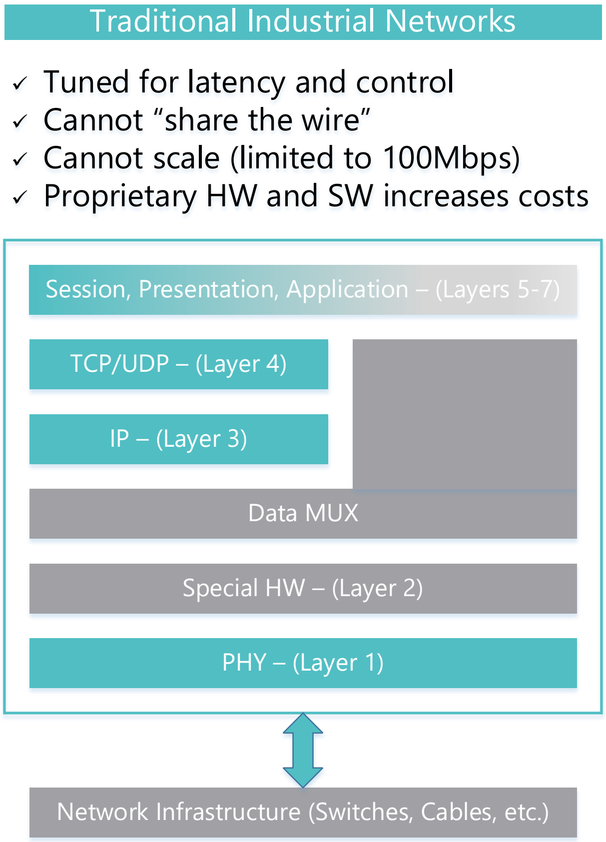

Figure 2 - Ethernet-based industrial networks are not compatible with one-another due to their proprietary layer 2 implementations, designed to provide the determinism applications require.

The challenge here was that, although Ethernet comes with a lot of standards, there was nothing that addresses the latency and determinism needs of industrial networks. While consumer devices can buffer video and audio to avoid disruption to playback of media, highly synchronised industrial processes can't afford to encounter such delays.

Many of the Generation I technologies developed new solutions that utilise an Ethernet PHY. Thus, PROFIBUS was complimented by PROFINET, CC-Link by CC-Link IE, Modbus by ModbusTCP, and DeviceNet by EtherNet/IP. Sercos, a hard-real-time technology using optical fibre, also developed Sercos III using an Ethernet physical layer. Further alternatives include EtherCAT and Ethernet POWERLINK. Despite using the same PHY, using proprietary layer 2 implementations means a mix of the above protocols still cannot be transmitted across the same cables, which does not tackle the installation complexity challenge.

Away from proprietary

All of these solutions needed to add proprietary elements to the typical Ethernet layers that weren't provided or ensured by the existing standards. These include ensuring latency could be guaranteed and determinism achieved, as well as making provision for reliability and fulfilment of safety integrity levels. While many commercial off-the-shelf (COTS) PHYs could be used, custom ASICs or FPGAs are needed to implement the augmented Ethernet hardware. Another challenge around Generation II Ethernet-based solutions is that they are limited to using 100 Mb/s PHYs, meaning that they won't scale to higher bandwidth needs.

Generation III of Industrial Automation networks are therefore coalescing around new standards in layer 2 that resolve these issues that, combined, are known as Time Sensitive Networking (TSN). Use of common standards means that COTS networking technology will be available that can be used by anyone, and development of this technology is also being driven by other application spaces, such as audio/video and automotive, and not just industrial networking. This work also ensures that higher bandwidths, such a gigabit Ethernet, can be used as well as single pair Ethernet (SPE) that allows wiring to be significantly simplified.

Some of the key TSN standards that deliver the synchronisation and latency benefits industrial networks need are:

- IEEE 802.1AS – Timing and Synchronisation for Time-Sensitive Applications – this mechanism shares synchronisation data between a grandmaster network node and all other nodes to ensure a common base reference clock. This is used to ensure a common, synchronous time base and is a profile of IEEE 1588.

- IEEE 802.1 Qbv – this standard provides further enhancements to ensure end-to-end latencies for applications by blocking low priority traffic during defined time windows. This is to support applications such as closed loop control over Ethernet through the use of a time-aware scheduler.

- IEEE 802.1Qbu – this standard defines the pre-emption methodologies in OSI model layer 2 that makes the Interspersing Express Traffic (IET) of IEEE 802.3br possible. This involves reducing the latency of certain traffic in a mixed traffic environment, such as by intercepting long, low priority traffic.

There is also ongoing work to draw together the relevant standards into use cases in the IEC/IEEE 60802 TSN Profile for Industrial Automation working group. This should ensure interoperability in deployment; select the relevant feature, options and default protocols for particular use cases; and provide guidelines for configuration. The ultimate goal is to have a single TSN Profile for Industrial Automation.

Tackling implementations of Generation III industrial networks and beyond

One choice for the implementation of TSN capable Ethernet hardware are FPGAs. However, the implementation of the TSN block alone consumes a large number of LUTs (LookUp Table) along with a lot of SRAM. Once combined with other functionality, developers are forced to consider large footprint devices that are power hungry and dissipate large amounts of heat.

Standard Network Interface Controller (NIC) solutions provide for a compact implementation that use very little power, and this is the direction that has been chosen with the TC9562 (figure 3). Featuring a PCIe interface, it can be connected to a larger system-on-chip SoC device in a programmable logic controller (PLC) or even integrated as a plug-in card within an industrial PC. The firmware necessary for operation is downloaded over PCIe or, if desired, custom application code can be embedded into an external serial flash. This enables implementers to use the integrated Arm Cortex-M3 processor inside the TC9562 to implement additional functionality or even acceleration functions.

The PHY interface is flexible, supporting SGMII, RGMII, RMII and MII, allowing 10Mb/s, 100Mb/s and 1000Mb/s interfaces to be implemented. The SPE T1 PHYs can also be used, supporting the growing trend to lighter and simpler wiring solutions.

Figure 3 - TC9562, operating in conjunction with a PCIe capable host processor, simplifies the implementation of TSN networks.

All of the necessary features of TSN are implemented, from the time synchronisation of IEEE 802.1AS to the time aware shaper (TAS) of IEEE 802.1Qbv with six queues available. The frame pre-emption function of IEEE 802.1Qbu together with IEEE 802.3br ensures that time critical packets can pre-empt the transmission of packets of non-time-critical data. Many of the standards are broadly complete, although some standards are still being fine-tuned. Thankfully, these features can be updated in firmware should they change. This makes the TC9562 the ideal candidate for integrating brownfield installations with greenfield TSN networks (figure 4).

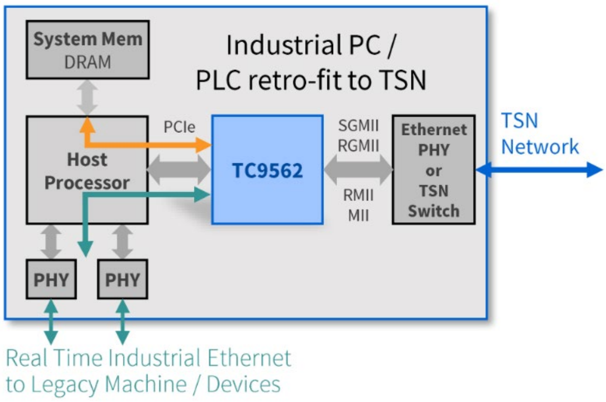

Figure 4 – Retrofitting of industrial PCs or PLCs enables brownfield adaption to greenfield TSN networks

A reference board is also available for testing and development purposes that can be used together with an off-the-shelf 64-bit industrial PC running Fedora (figure 5). The necessary drivers are available, and it functions together with a range of open source TSN stacks that include example applications required to build a TSN network. This is complemented by a reference software package from Toshiba that provides the necessary PCIe and eMAC drivers, firmware, sample PHY drivers for a range of silicon devices, flash loader, and other utilities and debug tools. It also includes bridge mode samples and TSN demos. In order to ensure developers always have access to the latest code and support, a Developer Zone has been created where developers can register to gain access.

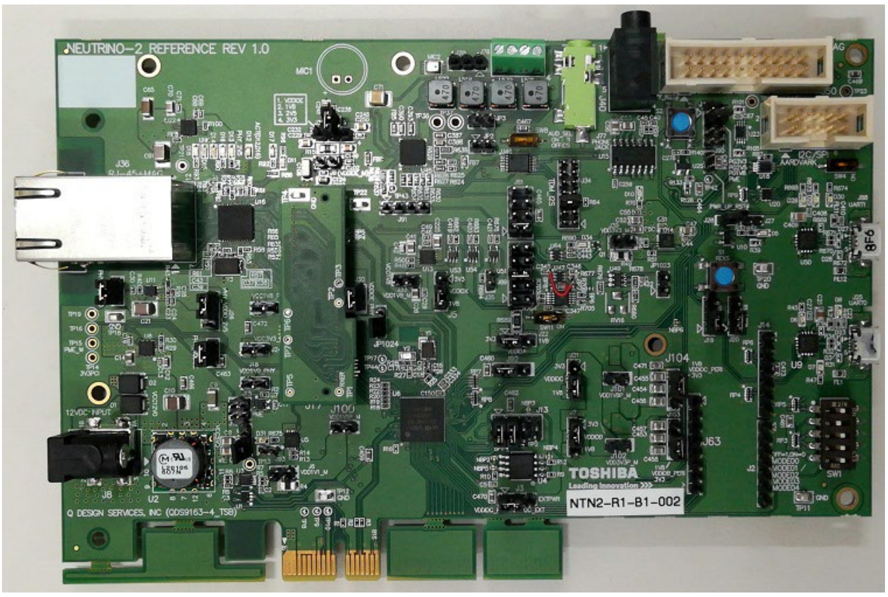

Figure 5 - The TC9562 reference board can be inserted into an available PCIe slot.

The TSN Demo allows the TAS function of IEEE 802.1Qbv and frame pre-emption capabilities to be visualised. Transmitting small data packets every 125 µs between a TSN Talker and a TSN Listener, impact on bitrate can be reviewed using tools such as iPerf[6]. The determinism of TAS can be switched on and off to review its impact (figure 6).

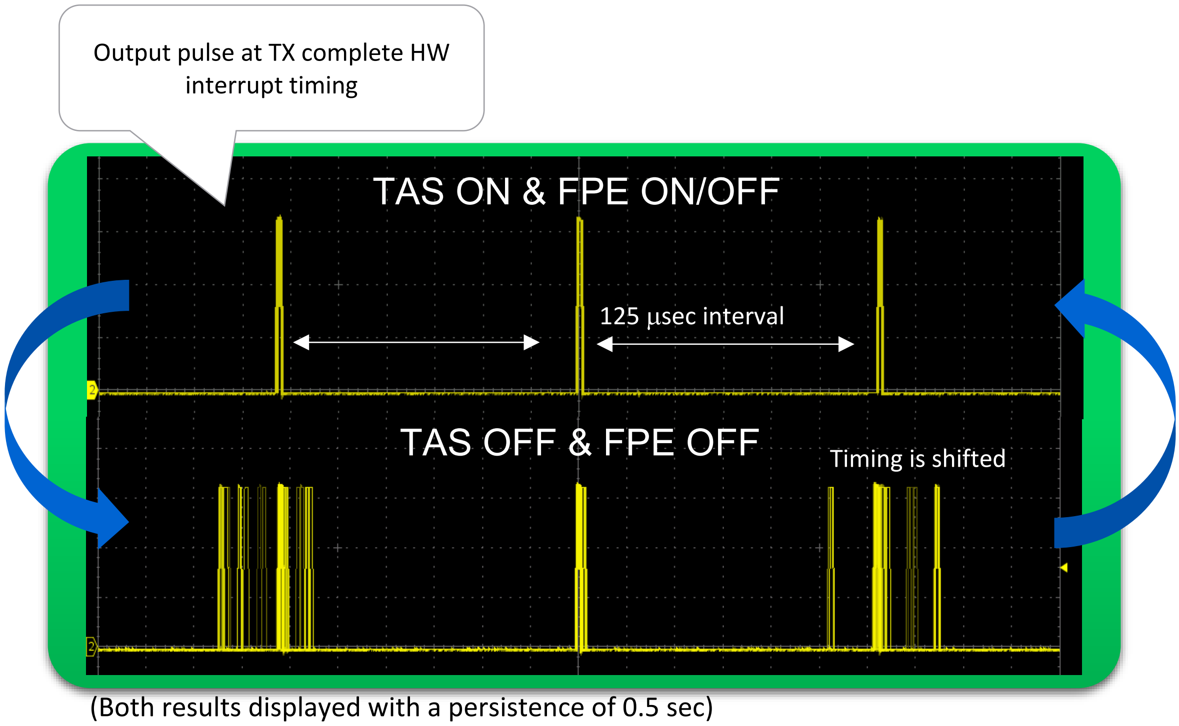

Figure 6: The Time Aware Shaper (TAS) of IEEE 802.1Qbv is shown to enable determinism in this demo included with the TC9562 reference board.

What to expect in Generation IV industrial networks

With standardisation occurring in the both the PHY and layer 2 only part of the challenge of industrial networks is being tackled. Much like the world of PCs, laptops and smartphones, we expect to be able to develop an application on top a networking stack and be able to simply connect devices together. With differing solutions above layer 2, this is currently not the case, even when the main OEM players settle on TSN.

The OPC Foundation has formed a working group to tackle the challenges around Field Level Communication (FLC), with all major OEMs participating in its development. The goal is to develop the OPC Unified Architecture (OPC UA) machine to machine communication protocol that targets the exacting needs of industrial automation[7], even going so far as defining functionality from sensor and actuator to cloud. By being open, cross-platform and service-oriented, along with the use of a robust security model, it has the potential to simplify much of the complexity today's industrial engineers are faced with, negating the need for endless protocol adapters. As well as supporting the traditional client/server communications model, it has also adopted publisher/subscriber (Pub/Sub) messaging. This means that installation of new equipment can be much simplified with standard equipment, such as a robot arm, publishing its degrees of freedom, maximum lift payload capability, and further specifications in a simple and standardised data structure that can be used by those network nodes that need it and can interface and control it.

Summary

Industrial automation has been slow to adopt and roll-out networking technology. The significant investment required along with the slow pace of change have something to do with this. However, with more and more applications making use of bandwidth hungry technology, such as vision systems, and a need to simplify installation costs and maintenance, changes are occurring that will deliver on these needs. Ethernet has started to establish itself as the de facto physical layer, but a disparate range of upper layer implementations has held back the dream of interoperability. This is because traditional Ethernet never had provision for guaranteeing latency or reserving bandwidth.

The introduction of the standards that make up TSN mean that COTS solutions can now be considered for Ethernet-based industrial networks, driving down the implementation cost of such systems. In addition, the work to standardise the upper layers with OPC UA ensures the momentum for this change is also there. Devices such as the TC9562, with their out-of-the-box support for standard open source software implementations provide an ideal basis for upgrading to TSN-capable networking technology.

References

1. https://en.wikipedia.org/wiki/Modbus
2. https://en.wikipedia.org/wiki/Profibus
3. https://en.wikipedia.org/wiki/CC-Link_Industrial_Networks
4. https://en.wikipedia.org/wiki/DeviceNet
5. https://en.wikipedia.org/wiki/CANopen
6. https://iperf.fr/
7. https://opcfoundation.org/wp-content/uploads/2018/11/OPCF-FLC-v2.pdf

Restrictions On Product Use

Toshiba Corporation and its subsidiaries and affiliates are collectively referred to as “TOSHIBA”.
Hardware, software and systems described in this document are collectively referred to as “Product”.

- TOSHIBA reserves the right to make changes to the information in this document and related Product without notice.
- This document and any information herein may not be reproduced without prior written permission from TOSHIBA. Even with TOSHIBA’s written permission, reproduction is permissible only if reproduction is without alteration/omission.
- Though TOSHIBA works continually to improve Product's quality and reliability, Product can malfunction or fail. Customers are responsible for complying with safety standards and for providing adequate designs and safeguards for their hardware, software and systems which minimize risk and avoid situations in which a malfunction or failure of Product could cause loss of human life, bodily injury or damage to property, including data loss or corruption. Before customers use the Product, create designs including the Product, or incorporate the Product into their own applications, customers must also refer to and comply with (a) the latest versions of all relevant TOSHIBA information, including without limitation, this document, the specifications, the data sheets and application notes for Product and the precautions and conditions set forth in the "TOSHIBA Semiconductor Reliability Handbook" and (b) the instructions for the application with which the Product will be used with or for. Customers are solely responsible for all aspects of their own product design or applications, including but not limited to (a) determining the appropriateness of the use of this Product in such design or applications; (b) evaluating and determining the applicability of any information contained in this document, or in charts, diagrams, programs, algorithms, sample application circuits, or any other referenced documents; and (c) validating all operating parameters for such designs and applications. TOSHIBA ASSUMES NO LIABILITY FOR CUSTOMERS' PRODUCT DESIGN OR APPLICATIONS.
- PRODUCT IS NEITHER INTENDED NOR WARRANTED FOR USE IN EQUIPMENTS OR SYSTEMS THAT REQUIRE EXTRAORDINARILY HIGH LEVELS OF QUALITY AND/OR RELIABILITY, AND/OR A MALFUNCTION OR FAILURE OF WHICH MAY CAUSE LOSS OF HUMAN LIFE, BODILY INJURY, SERIOUS PROPERTY DAMAGE AND/OR SERIOUS PUBLIC IMPACT ("UNINTENDED USE"). Except for specific applications as expressly stated in this document, Unintended Use includes, without limitation, equipment used in nuclear facilities, equipment used in the aerospace industry, lifesaving and/or life supporting medical equipment, equipment used for automobiles, trains, ships and other transportation, traffic signalling equipment, equipment used to control combustions or explosions, safety devices, elevators and escalators, and devices related to power plant. IF YOU USE PRODUCT FOR UNINTENDED USE, TOSHIBA ASSUMES NO LIABILITY FOR PRODUCT. For details, please contact your TOSHIBA sales representative or contact as via our website.
- Do not disassemble, analyze, reverse-engineer, alter, modify, translate or copy Product, whether in whole or in part.
- Product shall not be used for or incorporated into any products or systems whose manufacture, use, or sale is prohibited under any applicable laws or regulations.
- The information contained herein is presented only as guidance for Product use. No responsibility is assumed by TOSHIBA for any infringement of patents or any other intellectual property rights of third parties that may result from the use of Product. No license to any intellectual property right is granted by this document, whether express or implied, by estoppel or otherwise.
- ABSENT A WRITTEN SIGNED AGREEMENT, EXCEPT AS PROVIDED IN THE RELEVANT TERMS AND CONDITIONS OF SALE FOR PRODUCT, AND TO THE MAXIMUM EXTENT ALLOWABLE BY LAW, TOSHIBA (1) ASSUMES NO LIABILITY WHATSOEVER, INCLUDING WITHOUT LIMITATION, INDIRECT, CONSEQUENTIAL, SPECIAL, OR INCIDENTAL DAMAGES OR LOSS, INCLUDING WITHOUT LIMITATION, LOSS OF PROFITS, LOSS OF OPPORTUNITIES, BUSINESS INTERRUPTION AND LOSS OF DATA, AND (2) DISCLAIMS ANY AND ALL EXPRESS OR IMPLIED WARRANTIES AND CONDITIONS RELATED TO SALE, USE OF PRODUCT, OR INFORMATION, INCLUDING WARRANTIES OR CONDITIONS OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, ACCURACY OF INFORMATION, OR NONINFRINGEMENT.
- Do not use or otherwise make available Product or related software or technology for any military purposes, including without limitation, for the design, development, use, stockpiling or manufacturing of nuclear, chemical, or biological weapons or missile technology products (mass destruction weapons). Product and related software and technology may be controlled under the applicable export laws and regulations including, without limitation, the Japanese Foreign Exchange and Foreign Trade Law and the U.S. Export Administration Regulations. Export and re-export of Product or related software or technology are strictly prohibited except in compliance with all applicable export laws and regulations.
- Please contact your TOSHIBA sales representative for details as to environmental matters such as the RoHS compatibility of Product. Please use Product in compliance with all applicable laws and regulations that regulate the inclusion or use of controlled substances, including without limitation, the EU RoHS Directive. TOSHIBA ASSUMES NO LIABILITY FOR DAMAGES OR LOSSES OCCURRING AS A RESULT OF NONCOMPLIANCE WITH APPLICABLE LAWS AND REGULATIONS.