Extension MEDIA

Embedded Intel® Solutions

Spring 2017

www.congatec.us
6262 Ferris Square | San Diego CA 92121
Phone: 858-457-2600| sales-us@congatec.com

Intel IoT Solutions Alliance

conga-B7XD

First COM Express Type 7 full compliant module.

Server-On-Module

- Intel® Xeon® D CPUs with up to 16 cores & 24 MB cache
- 32x PCI Express lanes, 2x 10 GBit Ethernet
- Smallest server board ever - 125 x 95 mm²
- Optimized for demanding real time applications
- Ideal for rugged micro servers for industrial environment

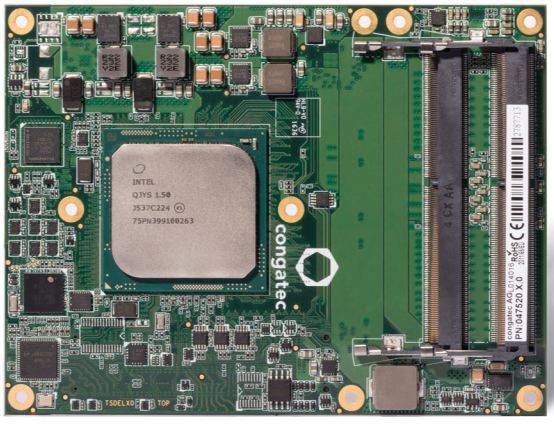

We simplify the use of embedded technology.

ADLINK
Leading EDGE COMPUTING

Industrial-IoT

For business-critical applications

COM Express Type 7 Module

Express-BD7

COM Express Basic Size Type 7 Module with Up to 16 cores Intel® Xeon D and Pentium® D SoC
(formerly codename: Broadwell-DE)

Compact/ Mini Size COM Express Modules

cExpress-AL / nanoX-AL

COM Express Compact Size Type 6 / Mini SizeType 10 Module with Intel® Atom™ E3900 series, Pentium®, and Celeron® SoC
(formerly codename: Apollo Lake)

Thin Mini-ITX Embedded Board

AmITX-AL-I

Thin Mini-ITX Embedded board with Intel® Atom™ E3900 series, Pentium®, and Celeron® SoC
(formerly codename: Apollo Lake)

SMARC® 2.0 - Smart Mobility Architecture

LEC-AL

SMARC® Short Size Module with Intel® Atom™ E3900 Series, Pentium® N4200 or Celeron® N3350 SoC
(formerly codename: Apollo Lake)

ADLINK Technology, Inc.
Toll Free: +1-800-966-5200
▶ *info@adlinktech.com* ▶ *www.adlinktech.com*

intel IoT Solutions Alliance Premier

embeddedworld2017
Exhibition&Conference
... it's a smarter world

Hall 1, Booth 1-540

Embedded/IoT Solutions

Connecting the Intelligent World from Devices to the Cloud

Long Life Cycle · High-Efficiency · Compact Form Factor · High Performance · Global Services

IoT Gateway Solutions

E100-8Q, E200-9B

Compact Embedded Server Appliances

SYS-5028D-TN4T

Network, Security Appliances

SYS-5018D-FN8T (Front I/O)

High Performance / Short Depth Solutions

SYS-1028R-WMRT

Cold Storage

SYS-5018A-AR12L

High Performance / IPC Solutions

SYS-6018R-TD (Rear I/O)

- Low Power Intel® Quark™, Intel® Atom™ Intel®, Core™ processor families, and High Performance Intel® Xeon® processors
- Standard Form Factor and High Performance Motherboards
- Optimized Short-Depth Industrial Rackmount Platforms
- Energy Efficient Titanium - Gold Level Power Supplies
- Fully Optimized SuperServers Ready to Deploy Solutions
- Remote Management by IPMI or Intel® AMT
- Worldwide Service with Extended Product Life Cycle Support
- Optimized for Embedded/IoT Applications





Intel Inside®. Powerful Productivity Outside.

TITANIUM

We Power the Cloud™

Learn more at www.supermicro.com/embedded

03_Q3_SM_USP_070616_MasterFile_EmbeddedIoT_Pg21

Intel: Surrounding Us with Technology, Making Life Better

More than just striving to be ahead of competitors is fueling Intel industrial automation, autonomous vehicle, AI, and other achievements.

By Lynnette Reese, Editor-in-Chief, Embedded Intel Solutions

The Spring 2017 edition of Embedded Intel® Solutions covers industrial IoT, automation, advancements in automotive, intelligent surveillance, and more in advanced technology. A striking fact is that Intel® has always had a reputation for excellence and is driving forward in several areas of technology with promise of market growth. Brian Krzanich, according to Reuters, said that Intel is "always paranoid about the competition, always driving. And you know that we live or die by the performance of our product."

Figure 1: February 2017: Ford Fusion Hybrid data collection cars at Intel® Corporation's Advanced Vehicle Lab in Chandler, Arizona. The cars are part of the Automated Driving Group and drive nearby streets collecting information to enrich "deep learning" models for future automated driving. (Credit: Tim Herman/Intel Corporation)

Beyond its motivation to make shareholders happy, Intel is finding inspiration from technology's potential to improve people's lives, create new markets and jobs, and provide creative solutions to real-world problems. It's a great time to be an engineer. Technology can extend lives, reduce misery, create opportunity, and outfit the disabled with tools for living.

Intel's Xeon® Phi™ family of chips, with up to 72 cores, enables high performance computing that makes possible computer modelling and simulation in supercomputers. Intel's Quark™ processors are the engines of IoT devices worldwide, with powerful computing in tiny packages. And driving will never be the same with Intel automated driving platforms supported by 5G connectivity, artificial intelligence, and cloud-based support. Intel rode the wave that was Moore's Law, consistently providing smaller and more powerful chips for decades. However, chips have reached a barrier and are seeing processing improvement moving on to multicore solutions as improvements in smaller, higher performance, and more power efficient chips become harder to reach in a cost-effective manner. As Moore's Law peters out, we will see Intel rise to the challenge by expanding into new technologies with all the ingenuity that was formerly released upon improving desktop performance. What's next?

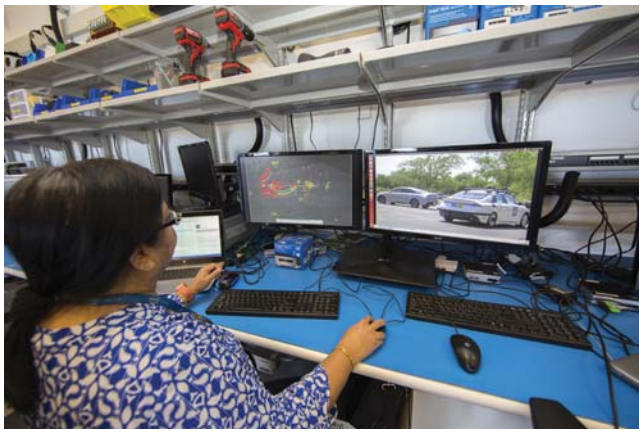

Figure 2: Intel's Nandini Sarkar, software lead for the Advanced Vehicle Lab, oversees data collection from cars at the lab in February 2017. (Credit: Tim Herman/Intel Corporation)

Embedded Intel® Solutions
www.embeddedintel.com

Vice President & Publisher
Clair Bright

Editorial

Editor-in-Chief
Lynnette Reese
lreese@extensionmedia.com

Managing Editor
Anne Fisher
afisher@extensionmedia.com

Senior Editors
Chris Ciufo
Caroline Hayes
David Bursky
Gabe Moretti

Creative/Production

Production Traffic Cordinator
LS Jerrett

Graphic Designers
Nicky Jacobson
Simone Bradley

Senior Web Developers
Slava Dotsenko
Mariam Moattari

Advertising / Reprint Sales

Vice President, Sales
Embedded Electronics Media Group
Clair Bright
cbright@extensionmedia.com
(415) 255-0390 ext. 15

Sales Manager
Elizabeth Thoma
ethoma@extensionmedia.com
(415) 255-0390 ext. 17

Marketing/Circulation
Jenna Johnson
jjohnson@extensionmedia.com

To Subscribe
www.embeddedintel.com

Extension MEDIA

Extension Media, LLC
Corporate Office

President and Publisher
Vince Ridley
vridley@extensionmedia.com
(415) 255-0390 ext. 18

Vice President & Publisher
Clair Bright
cbright@extensionmedia.com

Human Resources / Administration
Darla Rovetti

Special Thanks to Our Sponsors

ADLINK TECHNOLOGY INC. congatec iBASE Portwell SUPERMICRO WinSystems

Embedded Intel® Solutions is sent free to engineers and embedded developers in the U.S. and Canada who design with embedded Intel® processors.

Embedded Intel® Solutions is published by Extension Media LLC, 1786 18th Street, San Francisco, CA 94107.

EMPOWERING IIOT

embeddedworld2017
Exhibition&Conference

Rugged, reliable and resilient embedded computing solutions

WinSystems' embedded single board computers are designed to support a broad range of industry applications in challenging operational environments, from energy and transportation management, to industrial IoT and automation. Our industrial embedded computer solutions enable the collection, processing and transmission of real-time data requirements at the heart of your overall system.

From standard components to full custom solutions, WinSystems delivers world-class engineering, quality and unrivaled technical support. Our full line of embedded computers, I/O cards, and accessories help you design smarter projects offering faster time to market, improved reliability, durability and longer product life cycles.

Embed success in every application with ***The Embedded Systems Authority!***

EBC-C413
EBX-compatible SBC with Latest Generation Intel® Atom™ E3800 Series Processor

EPX-C414
Quad-Core Freescale i.MX 6Q Cortex A9 Industrial ARM® SBC

PX1-C415
PC/104 Form Factor SBC with PCIe/104™ OneBank™ expansion and latest generation Intel® Atom™ E3900 Series processor

SCADA | ENERGY | IOT | AUTOMATION | TRANSPORTATION

Single Board Computers | COM Express Solutions | Power Supplies | I/O Modules | Panel PCs

WinSystems®
The Embedded Systems Authority

817-274-7553 | www.winsystems.com
ASK ABOUT OUR PRODUCT EVALUATION!
715 Stadium Drive, Arlington, Texas 76011

intel IoT Solutions Alliance

QUALITY ISO 9001:2008 CERTIFIED

Embedded Intel® Solutions

SPRING 2017

DEPARTMENTS

SPECIAL FEATURES

PRODUCT SHOWCASES

On the Cover:
Solutions which are matched to the specific needs of the Industrial IoT demand both industrial process sector expertise and innovative cyber security approaches. (Image courtesy Intel Corporation)

Empowering the Connected World

Boosting IoT Designs from Edge to Cloud

- Intel® Core™ based module, IoT gateway, network appliance
- Energy-efficient embedded solution
- Industrial operating temperature
- 7+ product longevity
- High reliability

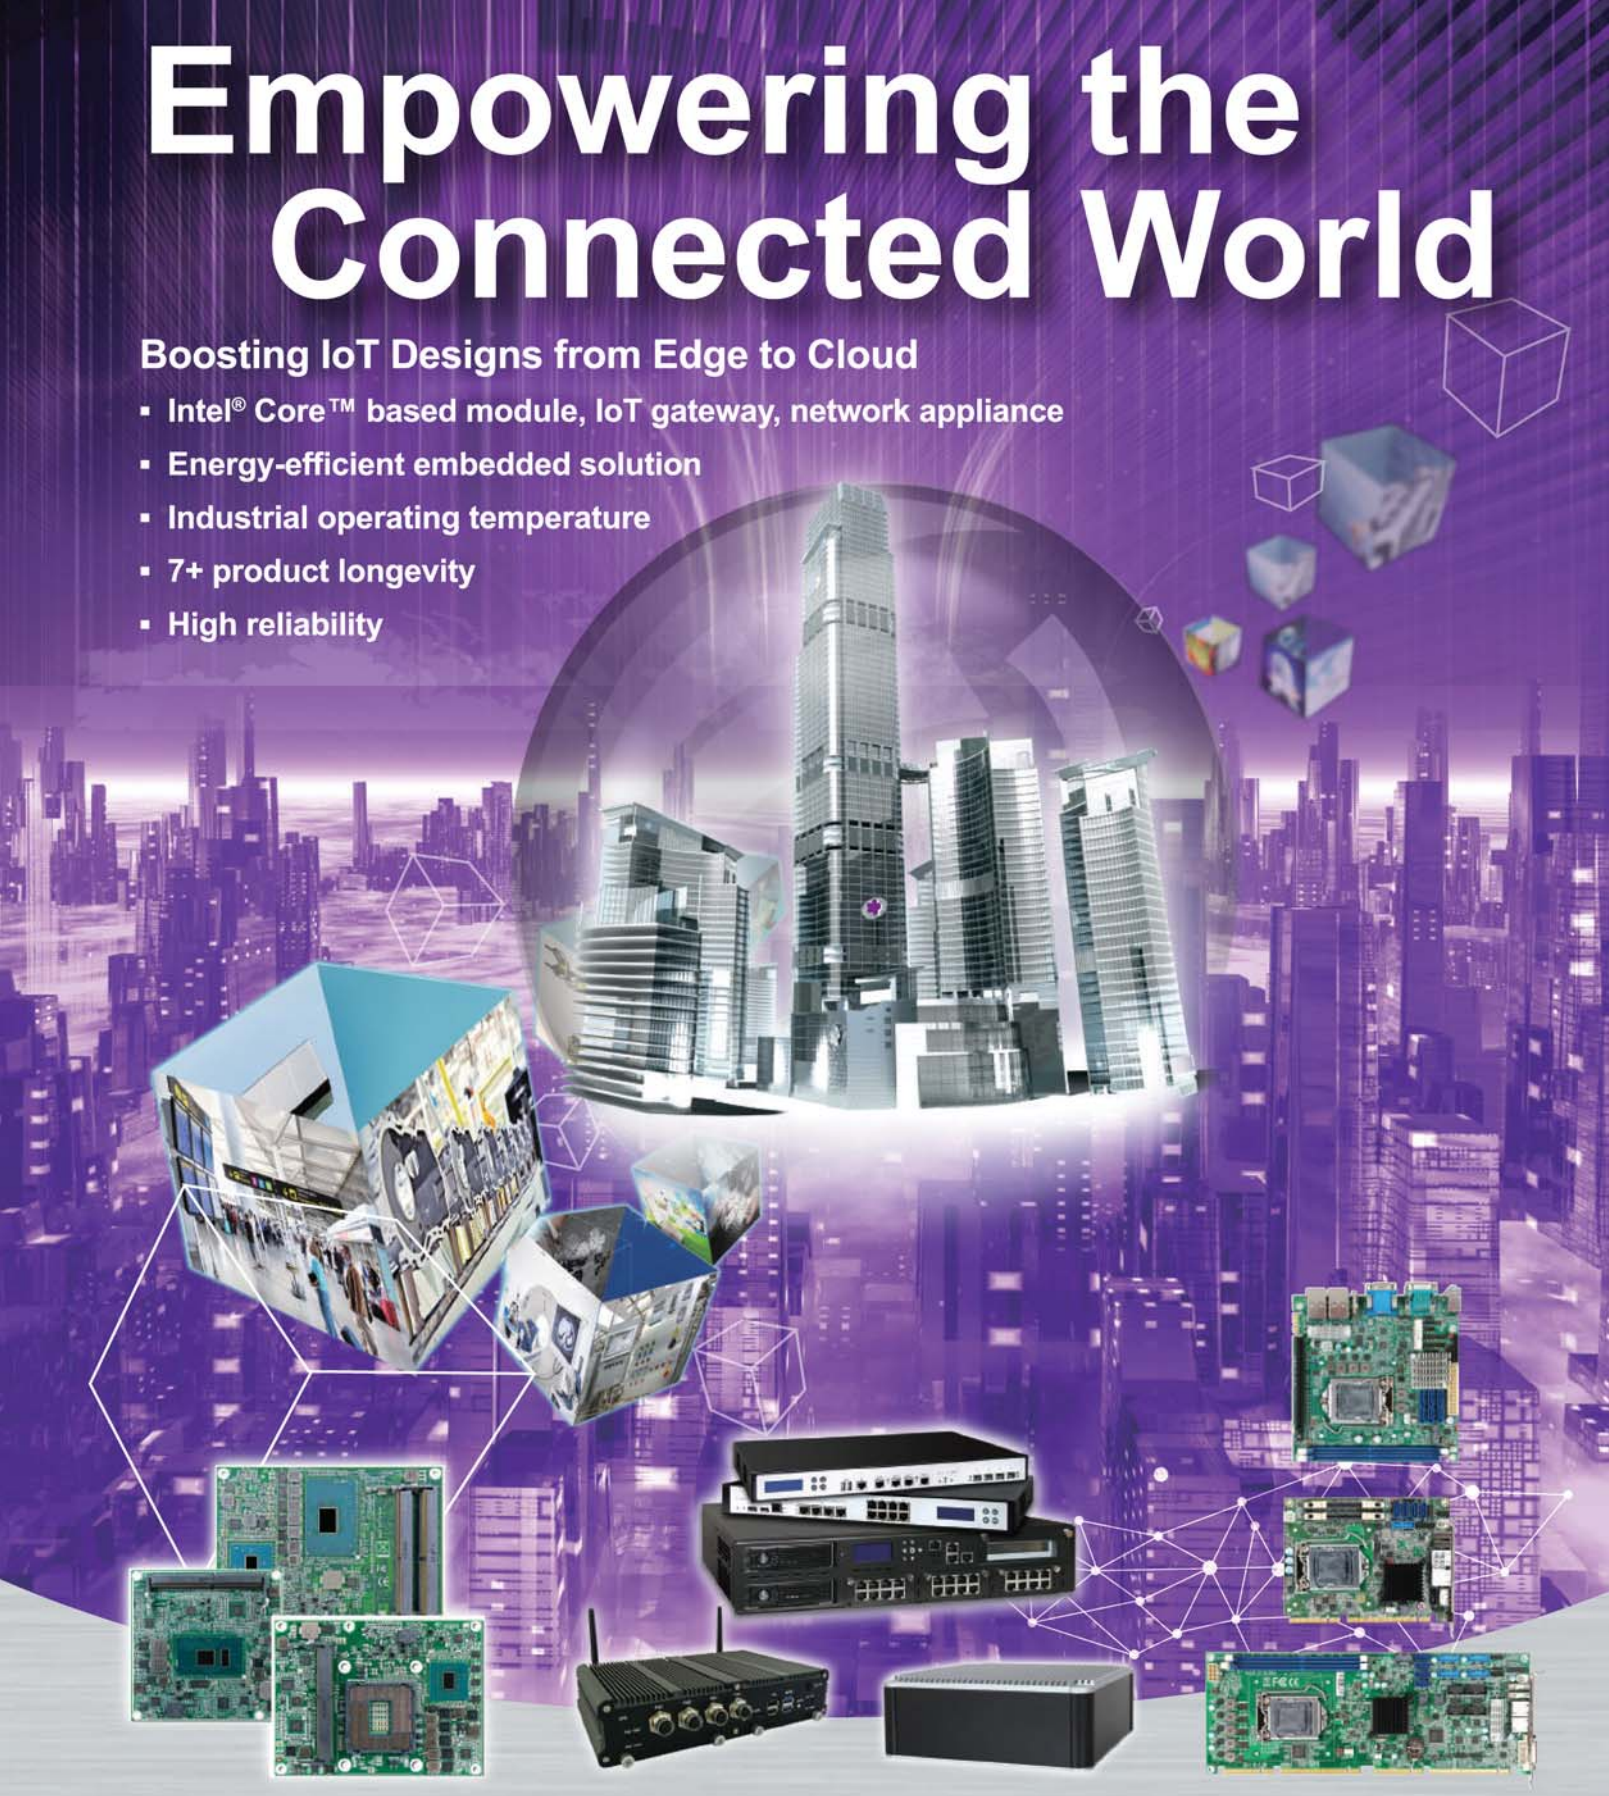

Enable Your Design Today

www.portwell.com
info@portwell.com
1-877-278-8899

Portwell

intel IoT Solutions Alliance

Consolidating Railway Communication

Virtualization can break down the barriers between various non-vital rail travel applications to cut costs and effort.

By Gunther Gräbner, MEN Micro Inc.

The number and variety of different automated services in modern rail travel has proliferated into a cost and management nightmare. Using Intel® Virtualization technology to bring these applications together into a unified environment can greatly increase efficiency and reliability, while reducing costs.

Modern railroad networks are a prime example of complex, rugged environments that require multiple computer systems for management, control, and convenience. These requirements go beyond the vital systems charged with controlling engines, fuel, braking, and more. Today's rail transportation also includes IT-type applications for energy management, passenger information systems, door control, ticketing, video surveillance, entertainment, and on-board Internet, to name a few.

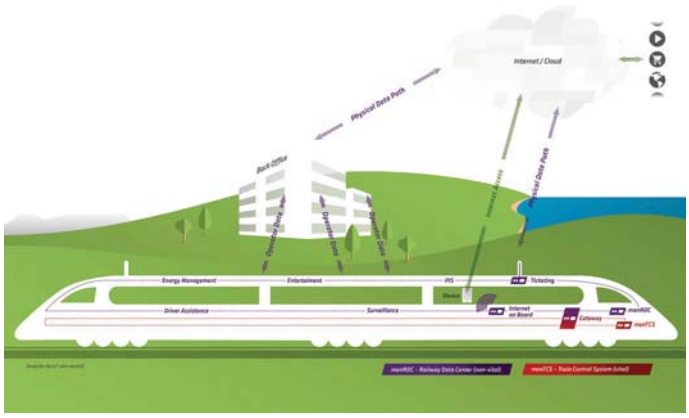

Figure 1: The rich variety of computer-based systems and services on modern trains has proliferated with a host of separate systems—Internet, video surveillance, ticketing, entertainment — that nevertheless must work together.

In addition, all these applications must communicate and exchange data at certain levels. Until now, this set of applications had been handled by separate and distinct computer systems. But the increased complexity of harnessing these isolated systems is growing more expensive and less manageable by the day.

Not only does each independent application have its own hardware, but most likely also requires a unique operating system and application software. This means railway management must deal with a host of different suppliers as well as communication and interoperability issues between systems. (Figure 1).

And it means difficulties with certification and obsolescence for each system, with some systems having far more capacity than they actually use, which translates to higher costs. For example, a ticketing system operates sporadically, and mostly while the train is in the station, so its dedicated computer sits idle the majority of the time.

If it were possible to bring these many disparate applications together into one computing environment—where they could share resources, exchange data more easily, and still fulfill their specialized functions—the savings in time and cost could be truly significant. To do this would require virtual computing environments, where each application could be transported with little or no alteration, and still run as if it were in its familiar old environment. Virtualization technology enables this scenario by networking virtual machines running different operating system environments and applications on a single, common hardware platform.

One such system family utilizing virtualization for just this purpose for modern railway computing is the menRDC. It combines different functions needed for an IT infrastructure on railway systems in a single, configurable and rugged package, providing a main server, storage system, and a network Ethernet switch. At the heart of the menRDC main server is a ruggedized CompactPCI Serial SBC with an Intel® Xeon® D-1500 (Broadwell) system-on-chip (SoC). This multicore hardware platform allows for the running of multiple virtual machines supporting different operating system environments. In addition, the SBC provides a rich combination of I/O interfaces including PCI Express, SATA/SAS, USB, and Ethernet plus signals needed for general system management. (Figure 2)

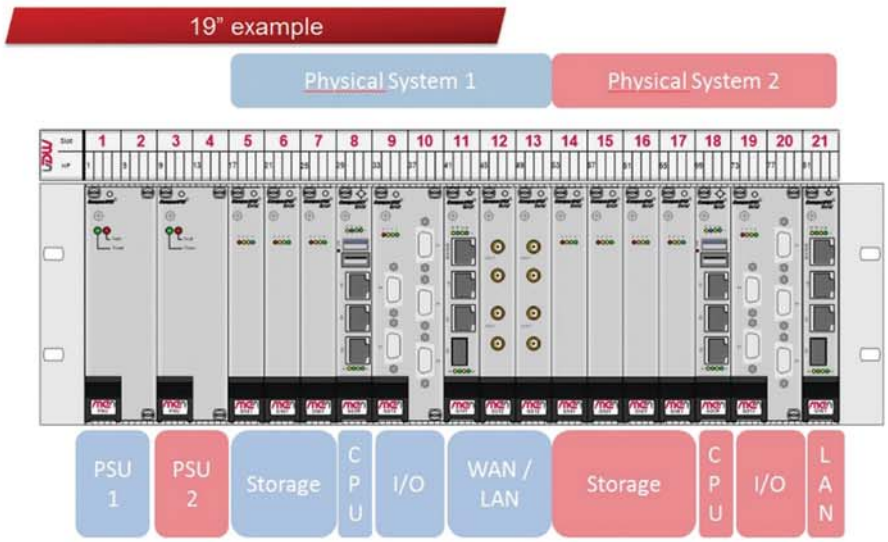

Figure 2: The menRDC consists of a variety of modules that are available in standard configurations or as built-to-order systems. Here are two complete systems in a single 19-inch rack.

The modular virtualization-based system is available in several pre-configured modules but can also be configured as a "built-to-order" system, thanks to the compatibility of the computing, networking, storage, and communication modules available.

For example, the main server, MEN Micro's MH70R, consists of one or two CPUs, each with a 16-core Intel Xeon D-1500, 2x10 Gb and 4x1 Gb Ethernet 3G/4G, Wi-Fi and GPS modem slots, up to four hot-swappable SATA slots and a pre-configured Linux operating system with drivers. Along with network switching and expanded storage units, these components can be combined in single- or multiple-rack systems with standard or custom configurations to meet a wide range of processing and capacity needs (Figure 3).

Virtualization Brings It Together

Given the power built into such a hardware design, how do we bring an array of applications together into this single platform, when many of the applications are running on different processors and operating systems? The answer is virtualization, which allows us to set up different execution environments, called virtual machines or VMs, on a single platform, sharing the underlying physical resources of that platform. The trick is to be able to do it efficiently, while maintaining the level of performance demanded by the various applications. The railway system described above does this with the combination of a

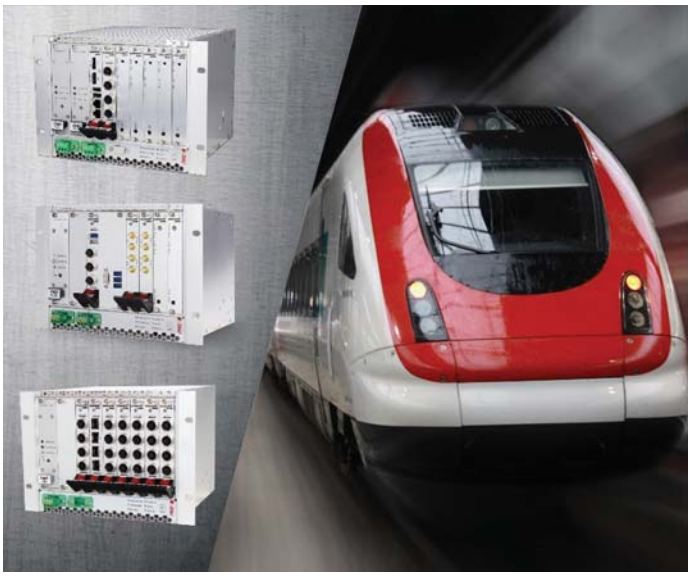

Figure 3: The menRDC pre-configured units include the main server, a storage extension, and network switch, while custom units can be made to order using a variety of available CompactPCI Serial boards.

lean and efficient type 2 hypervisor along with Intel's hardware assisted virtualization technology (Intel® VT).

The hypervisor abstracts the hardware from the application. It is hosted by Linux as a regular application and relies on the services of that operating system to manage system resources (CPU cores, memory, I/O, storage, etc.) for the applications it hosts in the various VMs it provides. Each VM has its own private virtualized hardware used by its application, as if it were on a separate processor environment. In reality, it's running on one or more of the D-1500's cores through the hypervisor. In addition to running the application, the hypervisor can be adapted to monitor other tasks, such as load balancing and crash protection. (Figure 4)

Any such virtualization involving a hypervisor is going to involve overhead, which will affect performance. Meeting performance demands of the applications is accomplished in two ways: First is the raw processing power of the D-1500 family, with eight cores on the D-1539 and 16 on the D-1577. The more cores, the greater the potential for workload consolidation. The other is Intel's hardware virtualization technology, which simplifies the software—and hence the overhead—used to enable and manage virtual machines.

The silicon virtualization technology addresses three areas. VT-x focuses on execution cores and CPU virtualization to reduce hypervisor complexity. VT-d focuses on direct memory access, remapping DMA transfers and interrupts for efficiency where the guest application is unaware of the

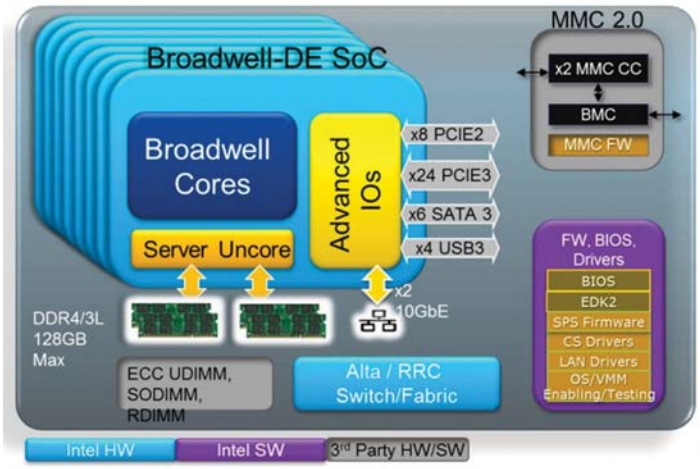

Figure 4: The Intel Xeon D family offers 4 to 16 Broadwell cores, which when used with a lean hypervisor and Intel's hardware virtualization technology, can provide the needed performance to incorporate applications from a range of different environments.

physical addresses. VT-c focuses on Ethernet connectivity, such that network devices are aware of VMs and will have Rx/Tx queues dedicated for each VM, which reduces a great deal of remapping without involving the hypervisor.

Optimized Railway Networks

The combination of rugged, modular and powerful multicore hardware with efficient virtualization technology in both software and hardware results in a single platform running many applications. Applications that formerly ran on different, dedicated hardware platforms are freed from hardware dependence and can share resources while running, as happens in the menRDC platform.

The savings to maintenance and upgradability are significant. If greater capacity is needed, the system can easily be increased by simply installing additional modules. The applications can efficiently share information as needed. Management can devise a user interface that conveniently accesses all needed data and control functions. Thanks to virtualization technology, modern railway systems can work on a common platform, reducing integration complexity and obsolescence issues, while better streamlining the needed system performance.

Gunther Gräbner has working in the automotive industry for more than eight years as a hardware developer, technical project leader and team leader of hardware development. Since 2011, he has supported MEN Mikro Elektronik GmbH as product manager for customer projects and MEN standard products.

HD Mobile Video Surveillance: Surpassing Network Limitations

Intel-based network appliances offer power densities that can help unburden your core networking infrastructure.

By James Piedra, Lanner

Many companies are just now starting to upgrade their video surveillance infrastructure to high definition (1 Mega Pixel +), and with the technology growing cheaper by the minute, many more will soon be upgrading to the latest and greatest technology. This jump in quality doesn't just require upgraded cameras, it also needs significant supporting network/storage infrastructure. The amount of IoT devices is set to grow exponentially over the following years, and video surveillance is by far one of the most bandwidth-hungry applications, with equally large storage requirements.

Figure 1: An HD (720p) video surveillance screenshot and an SD (480P) still.

Better Resolution = Better Surveillance

Figure 1, showing an HD (720p) video surveillance screenshot and an SD (480P) still, illustrates how a 4x increase in resolution brings a near-equivalent increase in detail. Given this, it's easy to understand why a high-resolution camera (2MP or 1080p) is ideal for implementing useful analytics like facial recognition. Once we reach 4K resolution the amount of detail will open up even more applications. Small text on nametags, hand gestures, and even lip-reading can become viable.

The increase in overall detail can help improve context, make it possible to discern smaller objects, and provide a solid evidence-admissible recording for liability reasons.

Requirements and Cost Comparison of a single Digital Surveillance Camera				
Quality	Bandwidth Usage (Mbps)	24h of video storage (1 GB = 8000Mb)	Hard Disk Drive 24h @ 3 cents per GB	Solid State Drive 24h @ 20 cents per GB
SD (480p)	2	100GB	$3	$20
HD (720p)	4	200GB	$6	$40
Full HD (1080p)	8	400GB	$12	$80
UHD /4k	25	1600GB	$48	$320

Table 1: Bandwidth requirements and storage costs of a single IP camera.

Struggling to Scale for HD Video

Just a few 4k video cameras are enough to saturate the bandwidth of modern Internet connections, and even enterprise cloud service providers are hard-pressed to make cloud-based video analytics available on streams greater than 480p(SD). Even with the impending upgrades to telecom infrastructure, centralized systems simply cannot be cost-effective in the face of hundreds/thousands of bandwidth-saturating devices.

Table 1 shows costs for a single day of storage, but keep in mind that recordings are typically stored for much longer than that. What's more, network and supporting infrastructure costs aren't even factored into this equation. These are the main reasons video-streaming giants like Netflix rely heavily on a more distributed model by using content delivery networks (CDNs). Add in the fact that in-vehicle and rolling stock surveillance is moving to Solid State Drives (for reliability reasons) while running on less capable wireless technologies, and things start getting expensive fast.

Network Video Recorders (NVRs) began with one simple function, essentially a next-gen DVR for IP-based video cameras. But its optimal position close to the video source has been increasingly acknowledged and exploited by industry players. Today's most sophisticated systems are essentially purpose-built network appliances. One example of such a system, running specialized Intel® x86 video surveillance software, is shown in Figure 2.

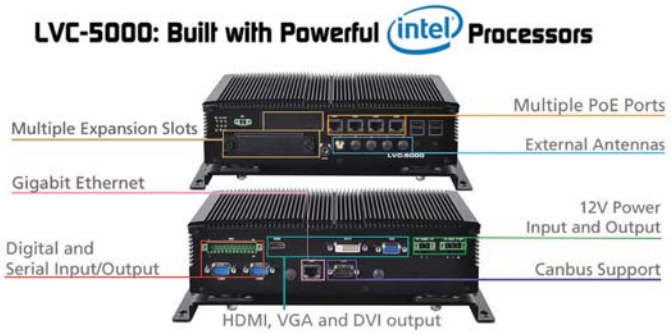

Figure 2: The ability to leverage a full-fledged ecosystem, encompassing both proprietary and open source software/libraries, make platforms which are Intel®- x86-based attractive.

X86-based appliances bring with them the major advantage of an extensive ecosystem of proprietary and open source software/libraries that make it the platform of choice in most cases.

The Secret Sauce: Powerful Analytics at the Edge

To avoid costly infrastructure and the embarrassment of DDoS'ing your own video surveillance infrastructure, you need to circumvent the bottleneck that is the round trip to the data center. To accomplish this, a growing trend is spanning across all industries which are tightly ingrained with information technology: Edge Computing.

By harnessing the great power-densities of Intel®-based network appliances and implementing efficient edge computing techniques combined with analytics, solution providers and even companies rolling their own systems have been creating highly scalable, cost-effective solutions. At the same time, powerful machine-learning algorithms at the edge are giving viability to many lucrative and beneficial applications.

Achieving Decentralized Analytics

Video-streams, even highly compressed ones, are by far the largest consumers of Internet bandwidth. There is no easy way around the heavy requirements of video-streaming aside from the incremental improvements encoders/decoders and specialized codecs like H.264/H.265 provide. But there is a way to strip useful information from all the noise in HD video streams and minimize bandwidth usage: Pre-process video streams as close to the source as possible.

In a world where Big Data continues to steamroll the opposition, metadata is king. Metadata is used by Google, for example, and is more or less the concentrated secret sauce of the data-driven businesses of today. Born from the need to efficiently classify, store, and transmit large amounts of information, metadata is data which describes another dataset.

Converting to Text-Based Metadata

Imagine you're law-enforcement and searching the feed for a suspect based on appearance. How would you go about it?

You could narrow down the search criteria by area (i.e., camera feed), but what about searching for people with a red hoodie? Or a license plate number? Certain hair color or facial features? An object travelling at xx speed? Such information is what an edge analytics-driven system would continuously transmit to the data center, instead of entire video feeds. Just as important, it would transmit unique and distinguishing data that is immediately useful for analytics. The techniques and algorithms involved in this process are highly sophisticated and still reaching maturity, as video feed information presents extremely noisy and unstructured data.

Storage Considerations

Modern Network video recorders house enough storage for several days, even weeks of HD video surveillance. This is enough to comply with law enforcement and more than enough for 90 percent of use cases. Video Management Systems (VMS) can intelligently manage recordings of importance in permanent storage based on analyzed metadata.

If cameras stream live feed to the operator room anyway, why not just centralize storage in that case? As seen in Table 1, there is a large difference in streaming SD, HD and 4k media. With the amount of bandwidth that's needed to transmit the feeds from a couple of HD cameras, you could easily stream many times that amount of low-quality feeds. While much lower in detail and definition, such feeds convey the benefit of allowing all streams to be viewed simultaneously. It's possible to look for the target and pinpoint the feed that needs to be accessed in full definition for further inspection. Now instead of needing core networking infrastructure capable of handling the full load of all the HD Cameras, you can effectively get away with 1/10th of the bandwidth. This technique is exploited in most modern NVRs, which all incorporate the capability to stream low-quality feeds and store high-quality versions locally.

More than Just Video Surveillance Moves Computing Closer to the Edge

IoT and mobile phones have ushered in a staggering amount of bandwidth-consuming devices, and exponentially more are on the way. Decentralized computing models are the only ones capable of keeping up with the ever-growing demands, and major players (Verizon, AT&T, the largest proponents of 5G and MEC) have already began morphing their networks in preparation for the upcoming challenges. We will be seeing less fallback/reliance on the cloud as businesses move to more scalable solutions.

James Piedra is a Network Platform Analyst at Lanner Electronics (lanner-america.com), a leading designer and manufacturer of Embedded Computers and Network Appliances. He researches and writes about the latest advances in Information technology, mainly focusing on Software-defined Networking, IoT, Cyber Security and mobile. On his free time James likes to tinker with consumer/embedded electronics and open source software (most time spent fixing something he broke in the process).

Intel GO Solutions Pave the Way for Autonomous Cars

If the destination is a place where solutions can take a big bite out of transportation costs…are we there yet?

By Lynnette Reese, Editor-in-Chief, Embedded Intel Solutions and Embedded Systems Engineering

In 1978, Cadillac introduced a trip computer, "a device set in the dashboard that is equivalent to a programmable calculator in the home. The computer can work out…a driver's fuel consumption or the number of miles before he reaches his destination." Car makers envisioned a driver information center that would extend dashboard controls, front radar scanners to determine a safe distance from the car in front, and keyless entry[1] (Marsh, 1979). Forty years later, we are promised self-driving vehicles by 2021, if Brian Krzanich is correct in estimating that this is when self-driving vehicles will be out and about with the Intel® GO™ system aboard.

SAE Level	Name	Narrative Definition	Execution of Steering and Acceleration/ Deceleration	Monitoring of Driving Environment	Fallback Performance of Dynamic Driving Task	System Capability (Driving Modes)
Human driver monitors the driving environment						
0	No Automation	The full-time performance by the human driver of all aspects of the dynamic driving task, even when enhanced by warning or intervention systems	Human Driver	Human Driver	Human Driver	N/A
1	Driver Assistance	The driving mode-specific execution by a driver assistance system of either steering or acceleration/deceleration using information about the driving environment and with the expectation that the human driver perform all remaining aspects of the dynamic driving task	Human Driver and System	Human Driver	Human Driver	Some Driving Modes
2	Partial Automation	The driving mode-specific execution by one or more driver assistance systems of both steering and acceleration/deceleration using information about the driving environment and with the expectation that the human driver perform all remaining aspects of the dynamic driving task	System	Human Driver	Human Driver	Some Driving Modes
Automated driving system ("system") monitors the driving environment						
3	Conditional Automation	The driving mode-specific performance by an automated driving system of all aspects of the dynamic driving task with the expectation that the human driver will respond appropriately to a request to intervene	System	System	Human Driver	Some Driving Modes
4	High Automation	The driving mode-specific performance by an automated driving system of the dynamic driving task, even if a human driver does not respond appropriately to a request to intervene	System	System	System	Some Driving Modes
5	Full Automation	The full-time performance by an automated driving system of all aspects of the dynamic driving task under all roadway and environmental conditions that can be managed by a human driver	System	System	System	All Driving Modes

Table 1: The five levels of vehicle automation as defined by the Society of Automotive Engineers (SAE) organization [Courtesy Intel]

Intel is correct to delve into automotive relationships concerning assisted and autonomous driving applications, since Level 5 autonomous cars are the holy grail not only for tech companies, but for agencies such as the National Highway Traffic Safety Administration (NHTSA). In 2014, the Society of Automotive Engineers (SAE) organization published a standard (J3016) for defining five levels of vehicle automation, from Level 0 at no automation through Level 5 at fully autonomous, with no human intervention. Level 1 is described as driver assistance, which might include cruise control, for example. Level 2 includes partial automation, where the system will execute aspects of both steering and acceleration and deceleration, for which adaptive cruise control qualifies. Level 3 is conditional automation, where the system not only executes Level 2 automation, but also monitors the driving environment. With Level 3 comes an expectation that the human driver will intervene if prompted to do so. One might consider Tesla's autopilot as a Level 3 task. Level 4 covers some modes of driving and continues all aspects of Level 3 but does not require driver intervention for that driving task. This may include a future driving mode for highway-only driving but not in hazardous weather, for example. Level 5 is fully

1. Marsh, P. (1979, December 6). The making of the computerised car. New Scientist, 770-773.
2. "Federal Automated Vehicles Policy." Department of Transportation. N.p., Sept. 2016. Web. 27 Jan. 2017.

Mark Rosekind, former head of the NHTSA (Source: allgov.com)

autonomous driving where all aspects of driving that humans would perform are done by the autonomous vehicle. No human intervention is expected at Level 5 (see Table 1).

Safety Is Just One Advantage

In 2016 there were 261.8 million registered cars and light trucks in the United States. According to former NHTSA Administrator Mark Rosekind, 94% of all vehicle accidents are due to human choice or error[2]. Highly Automated Vehicles (HAVs) have become a focal point for increasing safety in transportation. According to the Federal Automated Vehicles Policy, "While a human driver may repeat the same mistakes as millions before them, an HAV can benefit from the data and experience drawn from thousands of other vehicles on the road." The data from millions of automated and learning vehicles on the road will add safety to the list of benefits of big data. However, the NHTSA will "…continue to exercise its available regulatory authority over HAVs using its existing regulatory tools: interpretations, exemptions, notice-and-comment rulemaking, and defects and enforcement authority." Autonomous Vehicles (AVs) are a serious business and much more relevant than past innovations of convenience such as Cadillac's trip computer.

So, it's no wonder that fully automated driving is attractive; companies are racing to get autonomous cars on the road. Safety is just one advantage, however. Autonomous vehicles can ultimately reduce not only accidents, but traffic congestion, air pollution, and fuel consumption while reducing the need to expand highway infrastructure as self-driving cars change the American paradigm for mobility with Uber-like cost models for car ownership and use. Self-driving cars also offer independence to the disabled and the elderly. At a critical mass/point juncture, we will see a tipping point where AVs are the majority of the vehicles on the road with a potential for significantly bringing down the total cost of transportation. That tipping point, however, where a total number of AVs brings down the cost is a long way away; Rosekind also stated that older vehicles will remain in regular use on the roads an additional two or three decades after AVs come into serious use. Nevertheless, Rosekind's two-year tenure hallmarked the cooperation of the automotive industry in agreeing to make automatic emergency braking standard equipment by 2020, as well as a Federal Automated Vehicles Policy to "…ensure these technologies are safely introduced (i.e., do not introduce significant new safety risks), provide safety benefits today, and achieve their full safety potential in the future." Cars are safer than they used to be, but texting while driving can be added to the catalog that has listed driving under the influence for decades. Of late, technology that encourages driver distraction has increased recent highway death tolls in spite of technology-induced safety improvements.

The five largest chip makers currently serving the automotive market are NXP, Infineon, Renesas, STMicro, and Texas Instruments[3]. Intel wants to change that. At the 2017 Consumer Electronics Show (CES), Intel, BMW AG, and Mobileye announced at a joint press conference that their seven-month old partnership will soon produce 40 autonomous vehicle cars for testing on roadways by the second half of 2017. The companies have "developed a scalable architecture that can be adopted by other automotive developers and carmakers … from individual key integrated modules to a complete end-to-end solution providing a wide range of differentiated consumer experiences[4]." Mobileye, founded in 1999, is a leading supplier for core SoCs that go into vehicles with Advanced Collision Avoidance Systems. Mobileye's EyeQ® chip technology supports features such as vehicle and pedestrian "up ahead" warnings to support collision avoidance, all with a single camera. The partnership

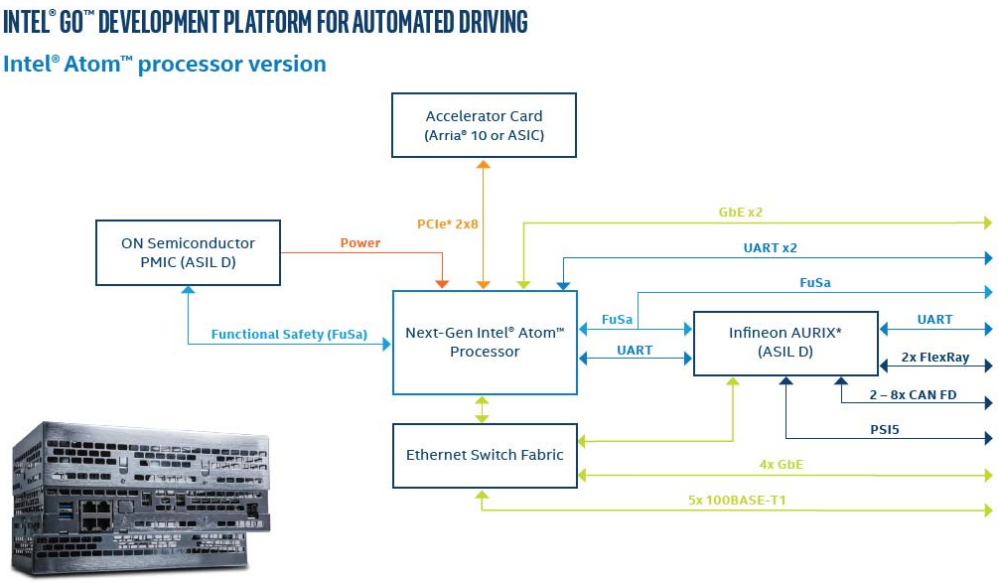

Figure 1: Block diagram of the Intel GO, the Intel Atom processor version. (Source: Intel.com)

3. WILMOT, STEPHEN . "How to Place Your Chips on Electric Cars." Wall Street Journal 19 Jan. 2017: n. pag. Print.

will extend Mobileye expertise to "the development of fusion algorithms … deployed on Intel computing platforms." Intel's computing power scales with solutions that include the Intel Atom™ or Intel Xeon® processors "with up to a total of 100 teraflops of power efficient performance without having to rewrite code[5]."

Intel's plans include not only high-performance computing for driver assistance to the autonomous car, but 5th generation cellular (5G) wireless connectivity and cloud. Intel GO Automated Driving solutions were introduced at CES 2017. The Intel GO In-Vehicle Development Platform for Automated Driving comes in an Intel Xeon version and an Intel Atom version, both with Intel Arria® 10 FPGAs for parallel processing. (Recall that Intel acquired FPGA-maker Altera in late 2015.) These boards provide a rapid and reliable way to develop, implement, test, and optimize everything from Advanced Driver Assistance Systems (ADAS) all the way through Level 5 fully automated driving without having to design hardware from scratch. The Intel Xeon version of the Intel GO platform enables solutions all the way to autonomous vehicles. The Intel GO automotive software development kit (SDK) allows developers to access the system for faster time-to-market with tools that incorporate computer vision and deep learning tool kits so they can develop and optimize algorithms for detection, sensor fusion, and execute on decisions. Sample reference applications for lane change assistance and object avoidance shorten the learning curve and time to market.

In addition to high-performance computing, Intel GO systems include software development tools, 5G-ready connectivity, a robust data center platform, and the latest in Artificial Intelligence (AI.) Automotive Intel Xeon processors, Intel Atom processors, and Intel Arria 10 FPGAs form the foundational basis for Intel's vision of autonomous vehicles. The included Intel Arria 10 FPGAs "feature hard floating-point digital signal processing (DSP) with speeds up to 1,500 giga floating-point operations per second (GFLOPS)." The Intel Atom and Intel Xeon-based GO platforms (with Arria 10 FPGAs) come with sample applications, run time libraries, and middleware.

Why 5G?

Intel supports connectivity in automated driving with the newly announced 5th generation (5G) cellular modem slated for release in the second half of 2017. 5G cellular communication is not going to be wide-spread until after 2020. Over-the-air (OTA) automotive updates allow car owners to avoid going to the dealership for updates. But 5G is more than just a channel for timely updates; it's a means for intelligent cars to communicate with other intelligent cars, with surrounding infrastructure such as smart "signs" with information on detours, speed limits, and warnings, and with pedestrians for a variety of reasons. Smart cars can communicate with a "smart" city to reduce traffic congestion and facilitate large gatherings such as conferences, marathons and festivals that require street closings. Presently, we can flip a turn signal to indicate intention, but 5G can be used to communicate a sudden need to swerve to avoid an obstacle. Thus, an intelligent car can alert other cars via 5G, as well as report roadway hazards to the local department of transportation. Not all 5G advocates feel that 5G is necessary for autonomous driving. However, the 5G Infrastructure Public Private Partnership (the 5G PPP), a European association that was initiated by the EU Commission, the telecommunications industry, small and medium enterprises, and researchers, believes that autonomous vehicles are not safe without 5G communication. Existing LTE cellular systems have latency that would negate a portion of the safety aspect that is gained with wireless communications to and from autonomous vehicles. To effectively support autonomous driving, wireless communication will need to meet minimum metrics for latency, reliability, throughput in heavy network traffic, and coverage[6].

5G will also allow fleet managers to monitor their fleet more closely, with more accurate knowledge of the location of vehicles and goods in a fleet of trucks, for instance. Fleet management with automation means having more accurate estimates for shipping arrival times, improved overall asset management, predictive maintenance and accurate maintenance records, and eventually eliminate the trucker's traditional role behind the wheel. Autonomous trucks will someday provide highly efficient driving control with lower instances of hard braking, softer starts, and constant attention to driving that will rival and surpass human driver capabilities. More efficient use of fuel and longer driving hours equate to improved profits. Autonomous cars will also make bus, taxi, and Uber drivers obsolete as scheduled and on-demand passenger pick-up and drop-off are usurped by self-driving vehicle services.

Lynnette Reese is Editor-in-Chief, Embedded Intel Solutions and Embedded Systems Engineering, and has been working in various roles as an electrical engineer for over two decades. She is interested in open source software and hardware, the maker movement, and in increasing the number of women working in STEM so she has a greater chance of talking about something other than football at the water cooler.

4. "BMW Group, Intel and Mobileye Will Have Autonomous Test Vehicles on the Roads by the Second Half of 2017." Intel Newsroom. N.p., 4 Jan. 2017. Web. 27 Jan. 2017.
5. "BMW Group, Intel and Mobileye Team Up to Bring Fully Autonomous Driving to Streets by 2021." Intel Newsroom. N.p., n.d. Web. 27 Jan. 2017.
6. Ma, Hui Sheng, Erqing Zhang, Shufang Li, Zhengnan Lv, and Jing Hu. "A V2X Design for 5G Network Based on Requirements of Autonomous Driving." SAE Technical Paper Series (2016): n. pag. Web. 27 Jan. 2017.

Ransomware and the IoT: Q&A with Brett Kelsey, VP and CTO for the Americas, Intel Security Group

Can we overcome our tendency to put function and availability ahead of security?

By Anne Fisher, Managing Editor

At the close of last year, which the December 2016 McAfee Threats Report suggested might be dubbed "The Year of Ransomware," Intel Security Group's VP and CTO for the Americas, Brett Kelsey, spoke with *Embedded Intel Solutions*. Kelsey offered his insights on one of the threats the Report highlighted, ransomware, its effects on automotive and other sectors, and related topics. Edited excerpts of our conversation follow.

Embedded Intel Solutions: What should the embedded engineering community know about one of the cyberattack species called out in the latest quarterly McAfee Threats Report, ransomware?

Brett Kelsey, VP and CTO for the Americas, Intel Security Group

Brett Kelsey, Intel Security Group: When you look ransomware's proliferation, including attacks we're starting to see more of, those based on proof-of-concept ransomware code, and what that means to the world of IoT, you start getting into things like the potential for ransomware to exist on the embedded systems and even to take over particular systems.

We demonstrated this at our FOCUS 16 Security Conference, where we took a car system and hacked it to show that we could put ransomware on it. Whether you are an ordinary driver or you happen to own a trucking fleet, or whether we're talking large business or individual person, ransomware attacks play out in a similar fashion: We've hacked your car and put ransomware on it, meaning you now must pay a certain amount of money—in the case of our demonstration at FOCUS it was one bitcoin, the equivalent of $700, to start your car back up.

Embedded Intel Solutions: How is ransomware affecting sectors in addition to automotive?

Kelsey, Intel Security Group: The medical field has been hit very hard by ransomware. Today, attacks on hospitals are going after their traditional IT systems. But their systems are inter-joined. The reporting component that you get out of a defibrillator or a heart monitor in some cases reports into a standard computer system, and if the standard system isn't functioning, you can't continue to perform healthcare.

Already, patient care has had to move from one hospital to another because of ransomware. And next generation ransomware is not going to focus 100 percent on attacking traditional IT; it's going after the true IoT space. We'll see more attacks directly into the IoT space.

The effects of such attacks could mean, for the Industrial IoT sector, that fluids at a factory might cease flowing through pipes that they should be traveling through, or start flowing through pipes where they do not belong. For the Smart Grid sector, hacking the systems that control power management could ultimately affect the grid itself.

Embedded Intel Solutions: What approaches to the ransomware problem you are describing will be most effective, and which less so?

Kelsey, Intel Security Group: There is a notional concept that we are going to be able to protect the devices themselves down to the device's chip layer, and we don't believe that is going to be effective.

There just isn't enough form factor that sits down at those chips. On the industrial side of the fence, you could have one particular chemical flowing through a pipe, and then that valve shuts off, and the next opens, and a different chemical is flowing through. And the valve itself, that's automated. It's IoT-controlled and has no idea what's going through it. In fact, it doesn't care. It just knows that it's either open or it's closed. And when it's open it needs to know and report how long it is open, how wide it opened, and when it's closed.

What must be baked into the chip is the ability to include attestation so that we can know what the device is doing, if it

Figure 1: Automotive is one of the sectors which can be harmed by ransomware attacks. [Courtesy Intel Corporation]

is doing it properly, and if it still is exactly the type of device that it is supposed to be.

Say for example, you are a coffee maker manufacturer. You want your IoT-capable coffee maker to act like a coffee maker, not like a different IoT-based "Thing." And so you need at the very least a reporting component that says: "Yes, I am what I say I am, and I continue to attest that I am valid and functioning the way I am supposed to be functioning."

That is about as much as you can expect at this moment from the IoT-related space, which means addressing more complicated security concerns requires adding some sort of an assistant. In our world that is cloud assistance, and now we are starting to see the cloud make large-scale analytics feasible. And that in turn will make it possible to provide the right level of feedback to the various IoT- or machine-based devices themselves. We're going to end up with that variety of connectivity—with the capability of going from the device itself straight to a cloud-based attached management and security platform, which will then do all the heavy lifting.

Embedded Intel Solutions: Please speak a bit about the intersection of the worlds of ransomware and machine learning.

Kelsey, Intel Security Group: From a developer perspective, steps can be taken now in the course of device development to put in the capability that will make it possible to start exercising future technology.

Machine learning as it exists today is still a fairly new thing. Several organizations have started on some form of capability, but it's still got a long way to go with regards to its true capability. That said, that should not slow down any developer from adding the hook components into the device manufacturing process they have today, and the code, so that any current and/or future security technology can be leveraged as soon as that capability comes on line.

Embedded Intel Solutions: What should the checklist be for the developer who needs to put in those hook elements?

Kelsey, Intel Security Group: At a minimum, the developers should be taking a look at all of the future function components that exist in their particular device or process that they are developing for. And as they are doing that, put the capability to provide attestation for each of those functions to say that they are (or are not) working as prescribed. That reporting capability can then be grabbed from a third-party solution of any kind, and then dealt with in the manner of which the third-party solution is capable.

Whether it's a medical device looking at various medical analysis components, a smart grid component looking at the power grids, or an industrial control monitoring gas or water flow, each one contains many functions. And each of those functions must be mapped out. What's also needed are capabilities to validate that those functions are working—that the device is functioning as normal—all the capability must be able to be grabbed from a management system that will continuously monitor the correct working state of each of the devices.

And then, at the same time, if you can take it to the next level (it depends upon room in the embedded component) you gain the ability to add true security protocols and encryption on the device side of the fence, so that your communication protocol can also be secure end-to-end. That would be the next layer that I would like to see for some of these IoT devices.

Embedded Intel Solutions: Is everyone on the same page with regard to understanding that solving ransomware attacks and other security problems means adopting a long-term strategy?

Kelsey, Intel Security Group: I am not sure that they are. And this is a cultural problem that we have from an industry perspective and even with respect to human nature. We tend to put function and availability way ahead of confidentiality and the security components that go with that. Organizations should take a really hard look from the bad guy's perspective—whatever it is they manufacture—and say, "What are the worst things that could happen in this scenario, and what would I ultimately do to provide capability so that I can stop them?

That long term pervasive view is not an easy thing to do. Yes, we have a consortium of car manufacturer and security protocols now that we have started to put together on the backs of various hacks of cars that have happened. Yet this is an example, despite the consortium's existence, that it takes a longer term of runway to get the devices themselves to, first, interoperate with each other in an overall system and second, have the security mechanisms baked in. You want those two things to happen so that you can protect those components from the onset rather than having an incident and having to do it after the fact. It's extremely difficult to do it after the fact.

Embedded Intel Solutions: A reactive mode is not want you want.

Kelsey, Intel Security Group: Especially in this area where you are dealing with things that touch so many people. And in a lot of cases, if you've got things embedded into a physical, hardware-based, functionality component, then the ability to change, modify, update and/or fix that becomes time consuming and extremely expensive.

Embedded Intel Solutions: Are there past IT security practices the industry should not abandon?

Kelsey, Intel Security Group: This gets down to the everchanging world of IT, and you are correct, there are practices that have been in place for quite some time and they still are very efficient and very effective. When you have the capability of putting a security agent on an actual device, and that agent has the capability of analyzing everything coming in and out of it and making decisions in real time on the device itself, that practice itself is extremely good and extremely efficient. We have gotten really good as an industry in how to deal with that.

The difficulty and the evolution is that the IoT world doesn't necessarily allow for that capability down at that compute component level. But given Moore's Law that everything gets twice as powerful with half the size form factor every 18 months to two years, that means that even the devices that are the smallest in size and nature will have more power and capability simply due to the sheer scale.

Sizes will get smaller and smaller. That will allow us—when we can—to put more capability down at the device level. And until that happens, we must do what we can to provide that capability outside the devices themselves. Until we can fully get there, we should take advantage of any opportunities to put security on the device as they occur.

Industrial IoT at Scale: What's Really Needed

Why cloud-centric architectures traditionally used in consumer IoT applications fall short when it comes to a larger class of IoT applications, especially those of the IIoT.

By Angelo Corsaro, Ph.D., ADLINK Technology, Inc.

Limitations of Cloud-Centric Architectures

Cloud-centric architectures are not applicable to a large class of IoT applications. Most notably, cloud-centric architectures fall short in supporting Industrial IoT (IIoT) systems and struggle with more demanding Consumer IoT applications.

The scary part is that the situation will only get worse with the predicted increase in the number of connected "things"—which could be anywhere from 50 -200 billion by 2020 according to Cisco, IDC and others.

But it's more than just the sheer number of "things" that is the problem. There is something more fundamental limiting cloud-centric architectures' applicability for IoT systems. Below, I've broken down these key fundamental issues.

Connectivity

Cloud-centric architectures assume that sufficient connectivity exists from the "things" to the cloud. This is necessary for (1) collecting the data from the edge, and (2) pushing insight or control actions from the cloud to the edge. Yet, connectivity is hard to guarantee for several IoT/IIoT applications, such as smart autonomous consumer and agricultural vehicles. As you can imagine, connectivity may be taken for granted in metropolitan areas, but not so much in rural areas.

Bandwidth

Cloud-centric architectures assume that sufficient bandwidth exists to bring the data from the edge into the data center. The challenge here is that several IIoT applications produce incredible volumes of data. For instance, a factory can easily produce a terabyte of data per day— and these numbers will only grow with the continued digitalization of factories.

Latency

Let's assume that the connectivity and bandwidth problem is solved. Is that sufficient? The short answer is no. There are still a large class of IIoT systems for which the latency required to send data to the cloud, make decisions and eventually send data toward the edge to act upon these decisions may be completely incompatible with the dynamics of the underlying system. A key difference between IT and IoT/IIoT is that the latter deals with physical entities. Reaction time cannot be arbitrary. It must be compatible with the dynamics of the physical entity or process with which the application interacts. Failing to react with the proper latency can lead to system instability, infrastructure damage, or even risk to human operators.

Cost

In the age of smartphones and very cheap data plans, most people assume that the cost of connectivity is negligible. The reality is quite different in IIoT due to either bandwidth requirements or connectivity points. While in consumer applications, the individual person—the consumer—pays for connectivity, in most IoT/IIoT applications, such as smart grids, it is the operator who pays the bill. As a result, the cost is usually carefully considered, as it affects OPEX and consequently operational costs and margins.

Security

Finally, even assuming that all the above listed issues are addressed, when security is an issue, a large class of Industrial IoT applications are either not comfortable with, or are prevented by regulations from pushing their data to a cloud.

In summary, unless you can guarantee that the connectivity, bandwidth, latency, cost and security requirements of your application are compatible with a cloud-centric architecture (1) you need a different paradigm, and (2) 99.9% of the IoT platforms available on the market are not of much use.

Fog Computing

Fog computing is emerging as the main paradigm to address the connectivity, bandwidth, latency, cost and security challenges imposed by cloud-centric architectures. The main idea behind fog computing is to provide elastic compute, storage and communication close to the "things" so that (1) data does not need to be sent all the way to the cloud, or at least not all data and not all the time, and (2) the infrastructure is designed from the ground up to deal with cyber-physical-systems (CPS) as opposed to IT systems. With fog computing, the infrastructure takes into account the constraints that interactions with the physical world impose: latency, determinism, load balancing, and fault-tolerance.

Software-Defined Automation, Digitalization and Fog Computing

As discussed earlier, cloud-centric architectures fall short in addressing a large class of IoT applications. These limitations have motivated the need for fog computing to address the connectivity, bandwidth, latency, cost and security challenges imposed by cloud-centric architectures.

Now let's consider some additional industry trends that are further motivating this paradigm shift and formulate a more precise definition of fog computing.

Two trends that are in some way at the core of the Industrial Internet of Things revolution are Software-Defined Automation, or Software-Defined Machines, and Digital Twins.

A trend that is disrupting several industries, Software-Defined Automation's raison d'être is the replacement of specialized hardware implementations, such as a Programmable Logic Controller (PLC) on an industrial floor, with software running in a virtualized environment.

Digital Twins, as the name hints, are a digital representation (computerized model) of a physical entity such as a compressor or a turbine, that is "animated" through the live data coming from the physical brother or sister. Digital Twins have several applications, including monitoring, diagnostics, and prognostics. Additionally, Digital Twins provide useful insights to R&D teams for improving next- generation designs as well as continuously ameliorating the fidelity of their models.

As Software-Defined Automation transforms specialized hardware into software it creates an opportunity for convergence and consolidation. Transform PLCs into software-defined PLCs, for instance, and suddenly they can be deployed on commodity hardware in a virtualized environment and decoupled from the I/O logic, which can remain closer to the source of data.

As a result of Software-Defined Automation and Digital Twins, there is an opportunity for modernizing the factory floor, consolidating its hardware, and increasing availability and productivity. Improved manageability, resilience to failure, and innovation agility also take place. Software-Defined Automation affords the opportunity to manage these systems as a data center. As this trend is influencing a large class of industries, it is worth highlighting that the transformations described above, along with the benefits, are not limited to industrial automation.

But there is a catch! The catch is that the large majority of these systems, whether in industrial transportation, or medical domains, are subject to the performance constraint already described. These systems interact with the physical world, so they must react at the pace the physical device imposes.

As a consequence, while traditional cloud infrastructures would be functionally perfect to address these use cases, they turn out to be inadequate as (1) they were not designed with these non-functional requirements in mind, and (2) they are often too heavyweight. Cloud infrastructures were designed for IT systems in which a delay in the response time may create a bored or upset customer, but will not cause a robot arm to smash against a wall or other machinery, or worse, hurt a human operator.

Fog computing[1] is not just about applying distributed computing to the edge. Fog Computing is about providing an infrastructure that—while virtualizing elastic compute, storage, and communication—also addresses the non-functional properties characteristic of these domains.

Fog computing makes it possible to provision and manage software-defined hardware, e.g., a Soft PLC, Digital Twins, analytics, and anything else that might be needed to run on the system while ensuring the proper non-functional requirements and delivering convergence, manageability, availability, agility, and efficiency improvement.

Deploying, monitoring, and managing software on the edge is made possible with fog computing's flexible infrastructure. Simply deploying some logic on an edge gateway isn't fog computing. Neither is fog computing traditional distributed computing.

Fog and Mist Computing

The attentive reader will have noticed that thus far we have focused on platforms that virtualize the compute, communication, and storage fabric available at the edge of the system. Yet, in many IoT systems, "things" have computational, communication, and storage capabilities that should also be exploited and managed uniformly. Thus, the natural question is, where does fog stop? Below the fog—on the devices—do we have something else?

Mist Computing

Mist is closer to the ground than fog, which in turn is closer to the ground than clouds. Mist computing is about bringing elastic compute, storage, and communication directly to "things". Thus, if we continue with the meteorological analogy, cloud infrastructure is high up in the data center, fog infrastructure is midway between the "things" and the cloud, and mist infrastructure is simply the "things."

Mist computing has two essential goals:

1. Enable resource harvesting by exploiting the computation, storage, and communication capabilities available on the "things."

2. Allowing arbitrary computations to be provisioned, deployed, managed, and monitored on the "things."

1. https://www.openfogconsortium.org/

As you can imagine, "things" in IoT applications are extremely heterogeneous with respect to platforms, resources, and connectivity. Thus, the main challenge for mist infrastructures is to be sufficiently lightweight to be able to establish a fabric that virtualize compute, storage, and communication without consuming too many resources.

Cloud, Fog and Mist Computing Convergence

If we take a step back and look from a distance at a generic IoT/IIoT, we will realize that from an infrastructural perspective we will have to deal with data centers that are in a public or private cloud, edge infrastructure, and the actual things. IoT/IIoT systems will need to exploit resources that span across these three tiers and provision, deploy, monitor, and manage applications and services across the three tiers. However, the landscape reveals complete fragmentation among the technologies used for cloud, fog, and mist computing. This fragmentation makes it hard to establish a unified end-to-end perspective on the system, and it makes it practically impossible to treat the system as a uniform and virtualized compute, storage, and communication fabric.

At this point the question is, "What can we do about it?"

The first step toward addressing a problem is recognizing it. To this end, the author of this paper has been raising the awareness around the challenges that this fragmentation may induce for the past year or so. The second step is to establish a vision of how the problem can be solved so that the industry can internalize it and eventually address it.

Let's focus for a moment on what would make sense for the user of an IoT/IIoT platform as opposed to the technical details of whether cloud, fog, or mist is the right answer.

From a high-level perspective, why should somebody designing an IoT/IIoT application care one iota whether he or she will be using cloud, fog, or mist computing paradigms? The only thing that really matters is that the platform provides a way to provision, manage, and monitor applications in such a way that applications can meet their end-to-end functional and non-functional requirements. The functional and non-functional requirements drive the allocation of application on "things," edge infrastructure, or cloud infrastructure, but anything else is just a detail, isn't it?

Fluid Computing

As a result, cloud, fog, and mist computing are now converging into fluid computing. Fluid computing is an architectural principle based on abstracting the topological details of the computational infrastructure. Fluid architectures provide an end-to-end fabric that can be used to seamlessly provision, deploy, manage, and monitor applications, regardless of whether the underlying resource is provided by the cloud infrastructure, the fog infrastructure, or by "things."

Fluid computing unifies under a single abstraction of cloud, fog and mist computing. Cloud, fog, and mist computing can be seen as applying fluid computing in a specific bounded context.

The impact of this line of thought can already be seen in the OpenFog Consortium Reference architecture, which now embraces some of the concepts of fluid architectures discussed above.

Making IoT Happen at Scale

In this paper, we have discussed the evolution of IoT architectures to support the expansion of IoT to more challenging and arguably more beneficial application domains, such as smart grids, smart factories, and autonomous vehicles. This is all good, but there is still something missing to really make IoT happen at scale—that is standardization. Today's reality is that IoT platforms are growing and continuing to fragment the market, interoperability is non-existent or extremely limited, and most IoT applications are silos with respect to connectivity.

To make IoT happen we need standards at a data exchange and data format level to be established. Some vertical applications seem to be standardizing over the DDS standard[2]. Others are standardizing over OPC-UA[3]. DDS tends to be preferred in systems required to operate at massive scale, with high performance and demanding fault-tolerance. OPC-UA, on the other hand, is widely used in the automation industry as a way of providing interactive access to field data. Both standard, along with defining mechanism for data sharing, provide mechanism for defining data models. DDS allows augmenting data models with QoS features that capture the non-functional requirements of the data. This is particularly useful for applications that need to control end-to-end QoS in order to ensure proper operations or quality of experience.

Standards exist that are ready to be adopted. End users need to be more aware of the importance of interoperability, and governments at a national and international level need to understand that, without interoperability, there won't be any IoT at scale—only a massive mess of stove-pipes clumsily integrated together.

Angelo Corsaro, Ph.D. is Chief Technology Officer (CTO) at ADLINK Technology Inc. As CTO he leads the Advanced Technology Office and looks after corporate technology strategy and innovation.

Earlier Corsaro served as PrismTech's (an ADLINK Company) CTO, where he directed the technology strategy innovation for the Vortex IIoT platform. He also served as Scientist at the SELEX-SI and FINMECCANICA Technology Directorate. There, he was responsible for the corporate middleware strategy, for strategic standardization, and R&D collaborations with top universities.

Corsaro is a well-known and cited expert in high-performance and large-scale distributed systems and with hundreds of publications in referred journals, conferences, workshops, and magazines.

2. http://www.omg.org/spec/DDS/1.4/
3. https://opcfoundation.org/

A Software-Defined Approach Sparks Digital Transformation of Industrial Automation

Can something be done to avoid the frustration of upgrades that arrive in the vendors' own sweet time and to overcome other obstacles?

By Gareth Noyes, Wind River

It's a pivotal time in the industry as many industrial companies are working towards/undergoing a digital transformation. Industrial companies and manufacturers have historically paid steep prices for automation systems purpose-built to perform a single task and lacking the flexibility to adapt to changing market environments. These proprietary solutions are not designed for interoperability with other products, which locks the buyer into the vendor and restricts component choice.

Working with a single industrial automation supplier may sometimes have benefits, but as technology advances and the marketplace demands and expects greater agility, the drawbacks become readily apparent. Proprietary systems are expensive to purchase (high CapEx) and maintain (high OpEx). Because these systems are developed in low volumes and built with highly specialized components, vendors lack the economies of scale inherent in commercial off-the-shelf (COTS) solutions.

Despite the Open Platform Communications (OPC) standard instituted in the 1990s, which enabled communications among proprietary systems, interoperability remains an issue. And as systems become increasingly interconnected, device and data security become paramount concerns. Security cannot be an afterthought. It needs to be designed in from the ground up, yet automation solution vendors often lack the experience to implement a layered security infrastructure leveraging multiple technologies.

A Cue from Telecom

The big issue, though, is that adding features and upgrading systems is costly and difficult, and usually takes place at the vendor's pace, constraining the operator from taking advantage of the latest technological advancements and innovations (Figure 1).

Industrial automation developers can take a cue from the experience of the telecom sector. At one time, telecommunications service providers also faced a predominance of proprietary equipment, which was a drag on the industry's growth. Over a dozen of the world's largest providers got together to lead the transition to interoperable solutions based on industry-standard servers—an approach called network function virtualization (NFV). After a few short years, telecom equipment vendors were able to offer software-based network functions running on COTS servers, making possible large economies of scale, wider vendor choice, and interoperability—all of which has benefited not only the service providers, but also the end users.

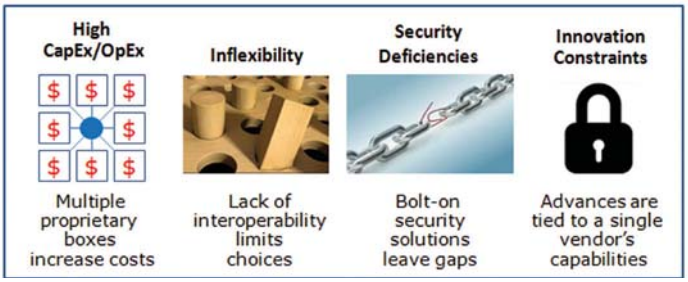

Figure 1: Roadblocks hindering Industrial Automation's evolution.

Now, a comparable digital transformation is underway in industrial automation, sparked by software-defined infrastructure and enabled by the IIoT. The premise of software-defined infrastructure is that most operations and control functions in an automation system can be consolidated onto standard, high-volume COTS servers capable of satisfying the real-time performance requirements of industrial environments. This creates an efficient, flexible and light-footprint alternative to proprietary industrial solutions. Software-defined infrastructure utilizes open standards and open platforms, extending them to meet industrial requirements, thereby reducing OpEx and CapEx and reaping the benefits of the IT cloud.

A software-defined infrastructure approach allows users, software vendors, and systems integrators to more easily develop interoperable components than proprietary solutions allow. Since software and server hardware are decoupled, software can be easily migrated and reused. Moreover, because the primary hardware platform is a server, it takes less effort to secure than custom platforms. The IT industry has developed various technologies for safeguarding servers that can be carried over to software-defined infrastructure to create robust and layered

security. And since software-defined infrastructure is based on open platforms, industrial companies are free to work with any supplier they choose to adopt the latest technologies and process innovations.

Flexible industrial automation, powered by a software-defined infrastructure, will enable companies to react more quickly and economically to an ever-evolving market landscape.

Higher Level of Interoperability

In contrast to conventional single-vendor, proprietary automation solutions, open software-defined infrastructure platforms allow for a higher level of interoperability among COTS components from multiple vendors, giving industrial users more choice and flexibility. An open platform approach also makes it easier to upgrade systems and add features to keep pace with changing market demands—all at a lower initial and ongoing maintenance cost than proprietary systems.

Let's take a closer look at how a software-defined infrastructure approach would work in an industrial automation application. The International Society of Automation's ISA-95 model, a standard for integrating enterprise and production control systems, comprises four levels as illustrated in Figure 2. Levels 1-3 represent the operations and control functions, while Level 4 is the enterprise-level business planning and logistics layer. The premise of software-defined infrastructure in automation is that most of the functions found in Levels 1-3 can be run on COTS servers capable of satisfying the real-time performance requirements of industrial environments.

More specifically, as illustrated in Figure 3, software-based digital controllers, PLCs/DCS, SCADA software, HMI, process historians, and applications in L1-L3 can run in an industrial software-defined infrastructure. Servers interface to sensors, actuators, and other physical industrial devices via distributed control nodes.

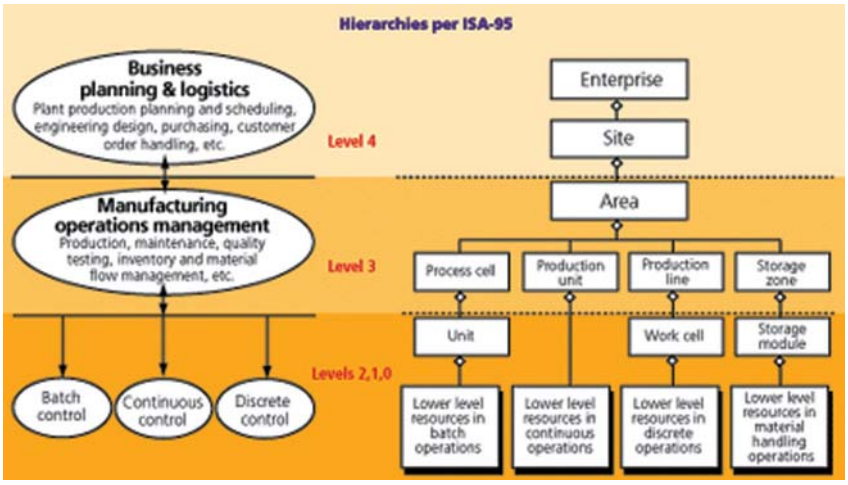

Figure 2: The four levels of enterprise and production control systems that are part of the International Society of Automation's ISA-95 model.

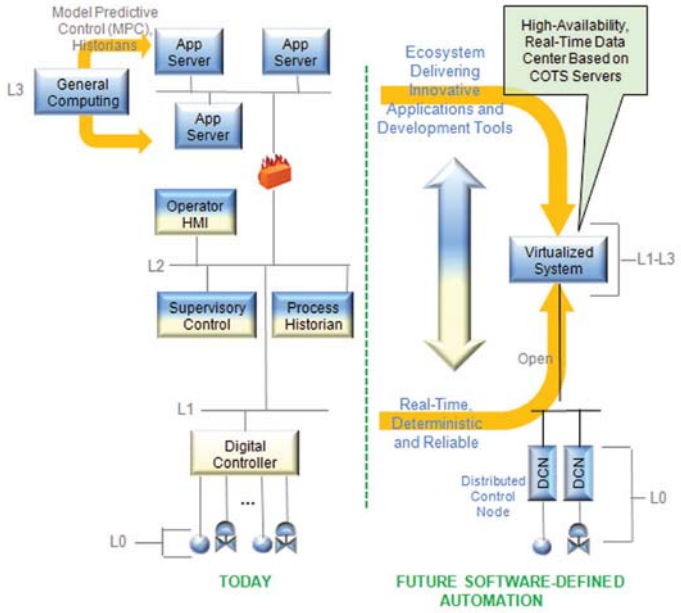

Figure 3: Satisfying the real-time performance requirements of industrial environments.

In contrast to a typical IT data center installation, software-defined infrastructure makes the data center "industrial grade," delivering the CapEx and OpEx benefits of an IT-based approach while satisfying such industrial requirements as high availability, real-time determinism, life cycle management, and hitless upgrades.

Compared to purpose-built proprietary solutions, software-defined infrastructure can deliver several cost-reducing benefits. It can lower hardware CapEx by substituting low-volume, custom computing platforms with a small set of high-volume COTS servers. These servers can be easier to manage than a sizable population of unique proprietary devices, lowering OpEx. Software-defined infrastructure would further lower CapEx and OpEx for logistics by significantly reducing the number of unique boxes that must be kept on hand for maintenance and the related costs to train and support staff on multiple unique articles.

A software-defined infrastructure approach has the potential to scale and expand with less effort because there are fewer wires, cables, and systems to deal with, minimizing connectivity related costs. Software-defined infrastructure solutions also take up less physical space near the industrial equipment they control.

By design, software-defined infrastructure -based systems require less field service support than traditional systems. In an industrial IoT environment, operators can monitor, diagnose and update software-defined infrastructure systems remotely and in real time, without deploying engineers, further reducing OpEx costs. If a failure occurs in the field, high-availability

system failover mechanisms help reduce the need for an emergency truck roll.

Finally, having a software-defined infrastructure can mitigate the costs of avoiding system obsolescence. The decoupling of functions implemented in software from the underlying hardware and software platforms makes it easier to update systems over their expected service lifetimes.

When you combine the cost benefits with the added flexibility and the ease of keeping pace with technological innovation, the case for software-defined infrastructure as the next wave in industrial automation becomes quite compelling.

Now, let's look at what is required in a software-defined infrastructure to realize those benefits.

Software-defined infrastructure automation solutions must run reliably and safely, gathering real-world industrial data and triggering real-time responses. To achieve this, a software-defined infrastructure must consolidate operations and control functions, and satisfy these criteria:

Low-latency virtualization: Software-defined infrastructure servers must support virtualization in order to run the diverse functions and applications found in industrial systems. The virtualization technology must have minimal overhead to realize real-time, deterministic performance for critical applications while optimizing resources for non-critical applications.

Deterministic networking: Fully deterministic, real-time communication via the IoT is essential for control functions in industrial environments. Time Sensitive Networking (TSN) achieves this by enabling a shared view of time and scheduling among industrial components.

High availability: In the event of software failure, software-defined infrastructure servers and applications must be able to perform automatic failover quickly enough to maintain control system integrity. Failover speeds need to be orders of magnitude faster than standard IT solutions. Carrier-grade telecommunication NFV solutions are approaching the automatic failover speeds needed for software-defined infrastructure. Virtualization technology facilitates failover in a number of ways—for example, restarting a clean backup software image without a reboot or turning control over to a full redundant server to avoid catastrophic failure.

Robust security: A software-defined infrastructure approach allows security technologies to be built in from the ground up across hardware platforms, middleware, applications, communications, and cloud infrastructure. The flexibility of software-defined infrastructure allows security solutions to adapt over time to respond to system and threat changes. Required technologies include secure boot, robust roots of trust (for example, Trusted Platform Module or TPM), digital random number generators, secure identities, local and remote attestation, anti-malware, data encryption, firewalls, authentication, authorization, and accounting (AAA), IDS/IPS, SIEM, and VPN tunneling.

Lifecycle management: Automation systems are typically expected to remain in continuous operation for years. Users must be able to perform lifecycle operations, such as software upgrades, live patching, capacity expansion, hardware updates and replacement, and physical and logical networking changes, without any loss of service. IoT solutions that allow easy installation, remote provisioning, and extensive monitoring of platforms, hardware, applications, and services are essential to maintaining system uptime and performance.

Enhanced platform awareness and monitoring: Software-defined infrastructure solutions need to support awareness of hardware and software status to guarantee required levels of service. IoT-powered platform awareness and monitoring capabilities enable automated resource allocation and reallocation as needed to adapt to change while maintaining performance, safety, and resiliency.

Best-in-class applications: Based on open x86 virtualization architecture using COTS hardware, software-defined infrastructure solutions must support the easy integration of IT technologies (Hadoop, Apache Storm, Java Analytics engines, Linux, and Linux containers). At the same time, solutions must implement operational technologies capable of satisfying real-time requirements (that are more stringent than IT) through the use of industrial strength real-time operating systems. System integrators and operators can then take advantage of the open platform to incorporate ISVs and best-in-class applications.

It all sounds fairly complex, but many of these infrastructure requirements have already been addressed by telecommunications networks that implement network function virtualization (NFV). One example is Wind River's fully integrated, full-featured virtualization software platform, Titanium Control. Designed to work with COTS hardware, it allows software-defined infrastructure automation solution providers to jumpstart development and focus on building applications rather than infrastructure. That, in turn, will enable industrial companies to accelerate the transformation toward software-defined infrastructure-powered automation and the many benefits it can deliver.

I predict this will be a revolutionary year with respect to digital transformation in industrial automation and the IIoT.

Gareth Noyes is Chief Strategy Officer responsible for overseeing Wind River corporate strategy and mergers and acquisitions activities, as well as for leading the Chief Technology Office. Charged with developing the company's long-term technology vision, he has Wind River positioned to address the evolving market landscape and disruptive forces such as the Internet of Things, the virtualization of the network, and the transition to a software-defined world.

Industrial Assets Often Outlive Connectivity Technology—What Are You Going to do About it?

Neither wireless nor wired connectivity options are immune to the march of time, but IoT developers can adopt strategies for prolonging asset life.

By Alex Romero, MultiTech

Editor's Note: In February MultiTech joined[1] the Board of Governors for the IoT M2M Council (IMC). Intel is one of the organization's Board-Member companies.

The Internet of Things is showing its age—and I don't mean its infancy. Although media hype might lead you to believe that the IoT is brand new or next, I'm here to tell you that in the industrial world, IoT is old hat. Public transportation systems have relied on remote connectivity since the early 20th century, factory equipment started to become widely connected in the 1960s, and who can forget the groundbreaking introduction of the first connected cars with the launch of General Motor's OnStar, more than 20 years ago.

Figure 1: A turnkey analog-to-Ethernet/wireless converter which emulates the traditional dial-up PSTN network, using integrated or external cellular modems.

Of course, connectivity technologies have advanced tremendously in that time, and those advances are only accelerating. Which presents a serious challenge to anyone trying to get a 10- to twenty-year life out of capital assets. The earliest cellular networks were shut down back in 2008 and 2009, causing manufacturers across many industries to scramble for a replacement to their Time Division Multiple Access (TDMA)-enabled systems to ensure continuous service.

Today in the U.S. connected devices are going dark due to AT&T's 2G sunset, a move likely to be followed by carriers around the world sooner rather than later. Even wired assets are not immune to the march of time. The copper wire back-

Figure 2: A cellular modem which is based on industry-standard open interfaces. The modem shown is built around a Telit xE910 cellular module, which features an Intel® XMM™ modem, comprised of an Intel® X-Gold™ baseband and an Intel® SMARTi™ transceiver.

1. http://www.prnewswire.com/news-releases/multitech-joins-iot-m2m-council-to-promote-flexible-solutions-for-industry-300401826.html

bone of the world's first telephony systems began rolling out more than 100 years ago and is still operational today. However, the cost to maintain those networks is on the rise, and telecom companies are actively encouraging users to transition to lower-cost technologies by increasing costs and halting service level agreements.

So, what's an IoT developer to do?

Retrofit

The simplest, if least sexy, answer is to retrofit existing equipment with an add-on connectivity device (Figure 1). Analog to digital converters can easily upgrade an asset designed to communicate using Plain Old Telephone Service, aka POTS, to newer Ethernet or wireless communications, allowing users to get a few more years out of their assets.

Moreover, a host of industrial wireless modems, routers and gateways are available to connect unconnected assets already in the field, whether vending machines, factory equipment, generators, utility meters or countless other industrial assets. By connecting existing assets, equipment owners can enable predictive maintenance and new services to generate enhanced value to their customers and improved efficiency for their operations. What's particularly nice about this "bolt-on" approach is that the modem itself is relatively affordable and can be easily replaced in the future.

Design for the Future

For new equipment where there is an opportunity to design connectivity in from the start, developers should think carefully about how long the asset will be expected to operate in the field versus how quickly the embedded connectivity protocols are likely to change.

Technology Selection

Some communications technologies change faster than others. Wired solutions tend to maintain a longer lifetime than their wireless counterparts, as evidenced by that 100-year-old POTS line that may still be running into your mother's house. Compare that to 4G-LTE, which was introduced in 2008 and has already undergone six releases since.

For the Internet of Things, ubiquity, application suitability and ease of deployment are also important factors. Clearly, wires cannot reach every asset we may wish to monitor, while certain wireless options may be easy to deploy, but not able to provide the bandwidth or reliability of wired communications.

As the developer of a new product, you have the luxury to specify how your device will communicate and to fine-tune that connectivity not only to meet the needs of today's application, but also to help mitigate the potential disruption future technological advancements cause.

We are witnessing the rollout of new fiber optic networks capable of Gigabit Ethernet speeds, which will ultimately form the backbone of the 5G cellular networks. It is the first such major wired infrastructure outlay since the advent of cable in the 1970s and is designed to outlast copper-based wiring. If your asset is fixed, it may be worthwhile to consider plugging in.

If, however, your asset requires mobility, there are a variety of wireless options. There is the cellular network, which is inherently driven by short life-cycle, large bandwidth usage and likely to change frequently; local-area options like WiFi which are becoming increasingly ubiquitous, but often have security barriers for roaming applications, and a variety of new low-power, wide-area wireless protocols which have been designed with the IoT in mind and whose proponents broadly proclaim will not undergo significant changes for extended periods. Which of these may be suitable to your application depends upon its requirements and should be carefully vetted for performance, range, link budget and power management before you commit.

Figure 3: Now available on the market are comprehensive portfolios of cellular connectivity products, including cellular routers which incorporate Intel processors, optimized for M2M (machine-to-machine) communications.

Footprint Compatibility

Another way to "future-proof" your design for new products is to select a component provider that offers footprint compatibility across a range of communications protocols (Figure 2). Such compatibility allows you to layout your board just once to accommodate evolving technologies by enabling an easy replacement of one technology for another. You also preserve the majority of your device design.

Prepare to Upgrade

We may be reaching the limits of physics when it comes to connectivity hardware. 2016 could well go down in history as marking the end of Moore's law. But progress continues in software. Products in the field can be readily improved upon through over-the-air updates for many years to come. That is, if ongoing software upgradability is taken into account during the initial design phase.

The savvy engineer designs in the appropriate two-way communications, embedded processing and memory to ensure improving upon deployed assets over time.

When designing a connected device, evaluate and incorporate device management, connectivity management and application management software to prolong the viable lifecycle of any assets you or your customers deploy (Figure 3).

Vendor Selection

Though the number of organizations claiming to have silver bullet solutions for the IoT is growing (perhaps even faster than IoT applications themselves), it is wise to choose suppliers with a track record of success and which offer a substantive variety of solutions to address current communications challenges and a roadmap to address future needs.

The Internet of Things, as stated at the onset of this article, is nothing new—and it pays to work with organizations that have been there from the beginning.

Alex Romero is a Strategic business professional with more than 10 years of experience delivering profitable results. He is passionate about business management, people, industrial engineering and the integration of systems thinking into practices that drive growth. He has strong analytical skills and enjoys building mutually inclusive relationships that foster collaboration. Romero is currently a Product Manager, managing cellular modems and routers at MultiTech. He holds a BS in Industrial and Systems Engineering from Tecnológico de Monterrey in Mexico and an MBA from the University of Saint Thomas.

A Slice of Pi and the IIoT's Appetite for Diversity

Intel® Core™ i7 processing and open standards are meeting the IIoT's varied I/O, graphics, and expansion demands when just dessert isn't enough.

By Justin Moll, DFI Tech

Solutions aimed at Internet of Things (IoT) applications keep on coming, with Raspberry Pi and other high-volume, basic feature, low-cost computers forming the basis for much of the market. But many applications—particularly Industrial IoT or IIoT—will require more diversity in I/O, graphics, and expansion. Versatility in I/O from open standard specifications like Mini-ITX, COM Express, and SBCs like 2.5-inch Pico-ITX can be a significant benefit for the industry.

Figure 1: Leveraging open-standard boards such as Mini-ITX, powerful multicore Intel® Core™ i7 processing can be utilized in small enclosures with a wealth of I/O & storage options. Larger versions with expansion slots for frame grabbers & controllers help drive advanced automation control systems.

A Slice of Pi

Yes, the Raspberry Pis of the computing world are simple and offer "good enough" performance at a very low cost. With their low-power processors (typically about 1.2 GHz), GPIO, USB, HDMI, GbE port, Bluetooth connection, etc., such computers cover the basics. So, you can have good graphics, solid data processing, communications interface, and the basics in I/O. The memory is quite low with basic RAM, but for the basic applications it is enough. For simple devices, these features cover the gamut of needs quite well in a small and light form factor.

However, it doesn't take long for the features of the simple computing boxes to fall short in a breadth of applications. For example, when sensing comes into play such as in machine vision, robotics, and many human-machine-interface (HMI) applications, typically more features are required. IIoT applications gather data from devices, provide feedback and communicate among the various factory floor systems. Data analysis allows managers to make process and manufacturing improvements to become more efficient. To achieve this, multicore processing with enough PCIe lanes is required. Also, the benefits of scalability, expansion, and a versatile feature-set for multiple applications are critically important.

I/O, Memory, and Expansion

Taking an open standard architecture, with its rich diversity of implementations, vast ecosystem of vendors, etc., is a key start. Even the very small 2.5-inch Pico-ITX SBC has all the features of the small computing kits we have mentioned, but has the added flexibility of I/O interfaces, expansion, and storage interfaces. There is also significantly more memory with the advantage of DDR3 performance.

When sensing is a key part of the system, typically more efficient and larger capacity storage is required than the basic boxes can provide. Interfaces to SATA/SAS storage give room to expand as system requirements grow. And extra MiniPCIe, PCIe, or other expansion slots make future growth possible while offering the flexibility that allows the user's systems (automation equipment, etc.) to be used in more applications. If, for example, a

Figure 2: For the factory floor, fanless systems that are dustproof and rugged for shock/vibration instill a higher degree of reliability and longevity.

factory using automation equipment expands operations, that factory would be able to expand the number of production lines without starting from scratch.

Offering versatility in the Graphics and I/O ports is also beneficial. Some applications utilize legacy equipment, which may vary greatly on a customer-by-customer basis. So, having VGA and DVI options in addition to HDMI is a nice benefit.

Mini-ITX Systems

Open standard, rich-computing-option Mini-ITX systems address SFF computing with benefits not available in Pi and other very low-cost systems, including:

- Performance—with more powerful processor options such as Intel Core i series, multicore, and more PCIe lanes for lots of functionality
- Greater memory and storage capacity—helping to facilitate the increased functionality
- Expandability with PCIe—this is important for many automation systems for the use of controller boards, frame grabbers, etc.

For example, the unit in Figure 1 has a Mini-ITX board inside providing 2.66 GHz (3.33 GHz Turbo mode) of power at 35 watts in an air-cooled system. The dual core processor has up to 16 PCIe lanes, allowing the unit plenty of lanes for a 2.5-inch storage bay, an optical drive, GbE, COM and LAN ports, plenty of other I/O and graphics ports, and 8GB of DDR3 memory. Mini-ITX is standard and ready to go without any hardware customization. Although this example embedded system is small, there are smaller versions with lower-power processors for fanless or completely sealed versions—a "plus" in many factory floor settings.

SFF Enclosures

Today there are powerful processors in Intel Atom™ designs that have low power usage and dissipation requirements. This gives us more bang for the buck and facilitates the use of fanless systems. Many of the IIoT applications will require small and light enclosures to house the electronics. By using open standard motherboards and SBCs, there can be several optimally sized enclosures with various performance levels and interfaces. In some versions, just a panel can be modified to accept different I/O options. With multiple processor options, on the same form factor board, the exact requirements of processing power, PCIe lanes, and other features such as memory and integrated graphics can be achieved.

A significant amount of the designs of IIoT will require industrial-grade components and ruggedization. This is especially true for robotics or in applications where the computing system will be transported or moved often. With a conduction-cooled clamshell, a SFF computer can provide a higher degree of ruggedization and increase reliability while foregoing one of the most common points of failure—the fan. Figure 2 shows a conduction-cooled sealed system with a standard 4-inch SBC motherboard inside. Since the enclosure does not have a fan, the processor selected is a quad core 1.91 GHz Atom processor that uses only 10W of power. This is easily cooled with cold plates in the sealed enclosure. Although even more compact than the fan-cooled system mentioned earlier, this unit still has the space and connectivity for plenty of I/O and dual GbE ports. The unit also has a 2.5-inch storage bay, an mSATA storage module via one of the 4 MiniPCIe expansion slots, a SIM card slot, and an optional microSD socket. Another benefit of the conduction-cooled clamshell approach is that it adds a level of ruggedization. This example unit meets MIL specifications for operating vibration and can withstand 3G of operating shock at half sine wave 11ms in 3 axes.

COM Express

COM Express is an attractive alternative as an easily upgradeable and customizable approach. A carrier motherboard is designed into the application, with pluggable COM Express mezzanines that plug into them. As requirements/performance levels advance over time, newer generation standardized COM Express modules can plug into the same motherboard. The various small form factor size options and high performance of the architecture is a benefit. There are, however, typically higher up-front development costs.

Open Source vs. Open Standard

The open source groups tend to focus on specific product designs, where even the Gerber files, schematics, and mechanical drawings are included. This lends itself to monochrome, commodity products with little differentiation. Open specification/open standard groups on the other hand define focus on common interfaces for interoperable products rather than finished products. Multiple vendors contribute to the base definitions and interfaces, but the implementation can vary greatly. Industrial Automation and IoT require a lot of I/O and processor differentiation, so the open standards/specification route is often greatly preferred. There are also benefits of scalability, multi-vendor interoperability, and a broad ecosystem of proven products by utilizing open standard architectures.

More for Less

In IIoT applications the multicore processing, expansion, and storage are critical for the sensing and data analysis requirements of those systems. Using open-standard motherboards and SBCs as a basis for small embedded computers provides scalable, multi-vendor options with a wealth of I/O possibilities, which is a key requirement for automation and many other applications.

Justin Moll is a dedicated consultant for DFI Tech.

Industrial Strength IoT

The immediacy of edge and fog computing brings connected devices closer to the data sources. Capitalizing on this proximity, WinSystems has incorporated the Intel Atom E3900 series in its latest series of industrial boards for Industrial Internet of Things (IIoT) applications.

By Caroline Hayes, Senior Editor

One of the problems encountered with the ever-increasing Internet of Things (IoT) is that more data and processing needs are being pushed to data centers. The increased traffic places demands on bandwidth. It is also creating quality issues. For example, data sent to a server for processing can suffer loss due to video compression and travel time. Processing data at the device, or the edge, is an advantage for many applications, from industrial to retail projects.

In October 2016, Intel introduced the latest Intel® Atom™ processor, the Atom E3900 series. It is designed specifically to support edge and fog computing. Edge computing is where processing and storage are the functions of the connected device, and fog computing is where processing and storage functions are performed by connected devices between the data source and the Cloud.

Figure 1: Machine vision benefits from reduced latency from edge computing.

Among its feature set is the ability to handle more sensors and tasks across a wide temperature range (-40 to +85°C) for industrial and other applications. It also increases computing power compared with the earlier, generation 3 processors by a factor of 1.7, to increase bandwidth and memory speeds.

The Intel Time Coordinated Computing (TCC) technology synchronizes the peripherals and networks of connected devices for determinism. By enabling one microsecond timing accuracy across a network, it also addresses latency issues in industrial applications, such as a robotics or control. Within the industrial sector, the processor also lends itself to predictive maintenance and remote management. Its graphics engine enhances 3D graphics and video capabilities for visual data identification and analysis as well as manufacturing inspection.

Formerly known as the Apollo Lake processor platform, the Intel Atom E3900 is built into a compact Flip Chip Ball Grid Array (FCBGA) and is based on the company's latest 14nm silicon technology. An automotive-qualified version, the Intel Atom A3900, will address in-vehicle applications, such as dashboard and vehicle-to-vehicle communications.

Industrial PC/104 Form Factor

One of the first companies to incorporate the Intel Atom E3900 series is WinSystems. It has produced the PX1-C415

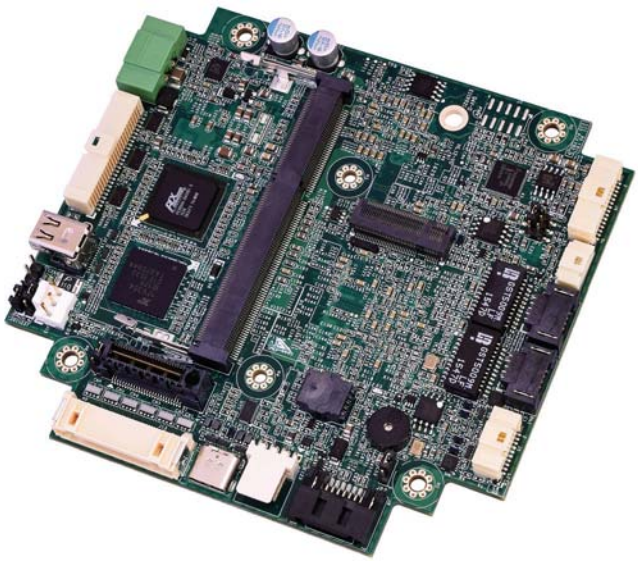

Figure 2: The PX1-C415 SBC maintains a compact form factor.

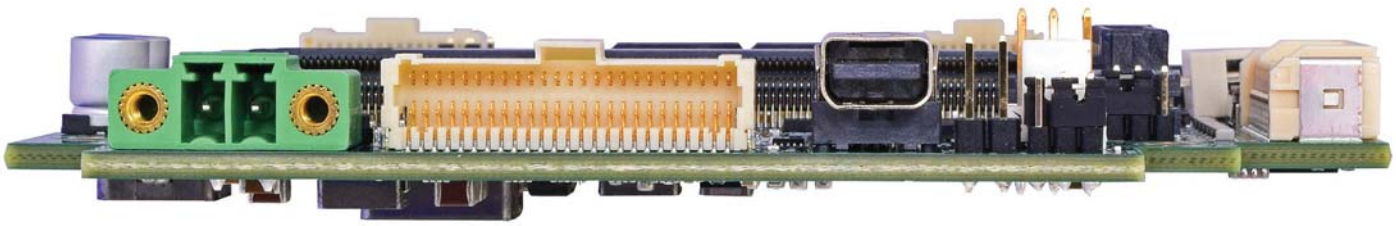

Figure 3: Expansion and connectivity options for the PX1-C415 include OneBank, USB 3.0 and M.2.

single board computers (SBCs), believed to be the first PCIe/104 OneBank single board computers based on the Intel Atom E3900 processor. In addition to the PCIe/104 OneBank expansion to support rugged applications, the SBCs can withstand temperatures in the range of -40 to +85°C and have dual Ethernet ports, dual video interfaces, four serial powers, 24 bi-directional general purpose input output (GPIO) lines and USB Type-C and M2 connectors. The SBCs support Windows 10 and Linux OS.

WinSystems designed the PX1-C415 to take advantage of the E3900's industrial temperature and ECC RAM support, says George Hilliard, Technical Sales Director, WinSystems. The design allows for mounting to an external heatsink when required and supports -40 to +85°C fanless operation. "Leveraging the USB 3.0 and DisplayPort performance, along with eMMC SSD and M.2 expansion, delivers a low-power, single board computer solution designed for the rugged, industrial environments our clients require," he adds.

Operating temperature is critical to the industrial environment, together with security. "The key features we considered when selecting the Intel E3900 series were the operating temperature range, ECC support, and improved Intel Security Engine," says Hilliard. "As security concerns continue for the plethora of IIoT applications, providing platforms with hardware encryption engines will become increasingly important. The low-power E3900 SoCs balance performance and power efficiency to support the latest operating systems such as Windows IoT Enterprise, Linux, and RTOS support."

Space to Differentiate

The PX1-C415's expansion options allow opportunities to expand and meet specific project requirements. In addition to the GPIO, the SBC provides Intel I210-IT Gigabit Ethernet interfaces, eight USB 2.0 channels, a SuperSpeed USB 3.0 channel, four serial COM channels, stereo audio and a watchdog timer.

Choice is important as IIoT requirements vary greatly, points out Hilliard. "Access to an ecosystem of data acquisition modules through USB 3.0, OneBank, and M.2 allows the PX1-C415 to be used in numerous applications," he adds.

Space is constrained in many industrial settings, so the size of embedded computing is an important feature. The PX1-C415 uses the PC/104 footprint with ever-increasing density, asserts Hilliard, to achieve dimensions of 4.55 x 4.28 inches (115.6 x 108.6mm). The company uses technologies in application processors, with expansion form factor to provide the required functionality that the IIoT demands, notes Hilliard. "Five years ago, a stack of four or five modules would be required for the same functionality now deployed on the PX1-C415 alone. Low profile expansion with OneBank and M.2 provide even more functionality in a very small overall package," he says.

Pre-installed OS and the ecosystem of expansion options and accessories free the engineer to focus efforts on the data acquisition and application software for the SBC. It is this which will differentiate the final embedded system solution, says Hilliard. "By providing expansion options such as PCIe/104 OneBank, M.2, and USB 3.0 Type C, the designer can maintain flexibility to add proprietary designs and be comfortable with future expansion options. It is a formula that has been proven successful for years and continues to be an advantage for small to mid-level production applications with minimal investment. Once an application has proven marketing viability, we often work with our clients to migrate to customized solutions utilizing the same or similar IP blocks to reduce costs and minimize any unused features."

Development samples of the PX1-C415 have been shipped and pre-production units will be available in Q1 2017. Full production is scheduled for early Q2, 2017.

Caroline Hayes has been a journalist covering the electronics sector for more than 20 years. She has worked on several European titles, reporting on a variety of industries, including communications, broadcast and automotive.

More Than Industrial Temperature at Stake: Q&A with Amit Gattani, Senior Director of Segment Marketing, Micron

A promising mint for the global economy, the Industrial IoT nevertheless needs the right approach to realize its potential.

By Anne Fisher, Managing Editor

Editor's note: Recently Embedded Intel Solutions spoke with Amit Gattani, Senior Director of Segment Marketing, Micron. He shared insights about mobile's effect on Industrial, leveraging memory's common denominator role to realize security benefits, and the effect an increasing reliance on supply chains is having, among other topics. Edited excerpts from our conversation follow.

Embedded Intel Solutions: What Industrial IoT features and trends should embedded engineers be keeping an eye on?

Amit Gattani, Senior Director of Segment Marketing, Micron

Amit Gattani, Micron: Certainly one feature to note is the Industrial IoT's potential to add up to US$14 trillion to the global economy by 2030[1]. Automotive and Industrial are areas that we are focusing on and we see substantial technical and business innovation going on in both those sectors. Industrial encompasses factories, transportation, surveillance, healthcare, and energy—for these and other sectors the business model change the Industrial IoT is bringing will generate significant operational savings as well as bring significant new economic factors into play.

A trend we are seeing might be called the "mobile waterfall effect." People in the Industrial IoT market, following in the mode of those in the mobile market, want quick, cheap productization, including leveraging mobile platforms. But while it's important to get something quick and fast, that may not be the best thing from a life cycle cost perspective or when considering total cost of ownership [TCO].

Especially where it concerns a device that becomes part of, say, a smart infrastructure like a surveillance camera or the parking/congestion management systems deployed in smart cities. The failure of any of these things can cause notable downsides to your business, much beyond the replacement cost of the hardware device. So, it is important to look at that aspect of it and not just ask, "How do I get the quickest device?"

Embedded Intel Solutions: How is Micron encouraging that attitude in the Industrial sector—one which considers TCO and not just, as with mobile, "How fast can we get to market?"

Gattani, Micron: Yes, we want a shift from this idea that if you could just translate a mobile or consumer-like platform into an industrial platform, then you could move things faster. That idea leads to thinking that picking a component with an industrial temperature, or 'IT' range and deciding, "Hey, if I get a wider temperature range product, I am going to be fine with it."

But the reality is—and this is why we have come up with this idea of Industrial Quotient or IQ—going beyond just IT—you need more than just an industrial-temperature product. You need ruggedization for vibration, shock, and thermal cycling. You need reliability that is not just one year but five to ten years. As an example, NAND technology is transitioning very rapidly from multi-level cells (MLC) to triple-level cells (TLC), and from planar to 3D. This is driven by insatiable demand for more flash storage capacity in mobile phones—everyone wants more storage in their iPhone.

As we move from MLC to TLC NAND, its program/erase cycle (P/E cycle) and endurance can go down a very steep curve. That may be acceptable for mobile or consumer applications because the life cycle of these products is meant to be just a couple of years or maybe three years at best. That endurance may not provide the five or 10 years' useful life needed for an industrial

1. Source: Accenture, Industrial Internet of Things: Unleashing the Potential of Connected Products & Services, March 2015

application. You have to really understand the reliability of the part you are selecting, not just at the initial stage of its life but for the latter part of its life cycle as well.

Embedded Intel Solutions: How does the Industrial Quotient approach come in part from customer conversations?

Gattani, Micron: The customers who have been in the Industrial market do understand why an Industrial Quotient approach makes sense, but these customers are also relying more and more on their supply chains. A GE, Tyco or Honeywell is not building a lot of the hardware itself. Their monetization is going to come from software, services, analytics and data bases. These companies are typically spending over 10 dollars on software for every dollar they spend building the hardware. The investment in software is so much higher now.

With the Industrial market's reliance on a fairly broad set of supply chains across the world, we have to get the 'Industrial Quotient' message to these supply chains and not just to the traditional customer base.

The role of the supply chains had the spotlight during last October's huge DDoS attack. Common to the surveillance cameras from a number of different OEMs, which the hackers used to create the attack, was a camera from Chinese firm Xiongmai.

So many people now come together to make a product, it's important to make sure everybody in your supply chain knows how to do the right thing and not just thinking it will be okay if the final end customer knows. Whether in the case of security-specific features or reliability-specific features, we have multiple examples of instances where such issues lead to field issues, which in turn ultimately increase the life cycle costs of the products. Obsolescence management is another major challenge that needs to be managed throughout the supply chain to drive down the life cycle costs.

Embedded Intel Solutions: One challenge with the Industrial IoT is that it can't be treated as a single entity.

Gattani, Micron: Yes, Industrial IoT applications are varied. One answer doesn't work for everything. There are different: system-level requirements; usage models; applications; read-write workloads; number of software updates you might get over the life cycle of your product. We understand end applications, including security demands, and the security features on our flash devices can be leveraged to build system-level security features.

And because the IoT, not just the Industrial IoT, is so fragmented, when it comes to boot security and being able to create the root of trust in hardware, flash memory is one of the best places to create that root of trust in hardware because it is the common element that goes across a lot of these devices.

On the other hand, the MCUs or CPUs for different applications are not the same and have a different capability and approach to implementing security. So, one of the threads of Micron's emphasis on Industrial Quotient Matters is to push the security down to the least common denominator, memory, because if you can address that in one place it applies to everything you build, irrespective of what MCU, CPU, or operating system you eventually use.

For example, in an automotive use case, you may have 40 to 50 ECUs in the car, and each ECU could [be based on] a different SoC architecture, but memory is the common thing in many of those. Wherever you can leverage the lowest common denominator to address a system-level problem, [that's the approach to take].

Embedded Intel Solutions: Do customers appreciate that given Micron's experience in a variety of sectors, there is opportunity for synergy and cross pollination?

Gattani, Micron: Absolutely. The semiconductor business requires very significant manufacturing investments and a lot of technology innovation and product development innovation to ensure we are providing the most leading edge products at best possible cost to our customers. If you are focused only on a smaller part of the market for industrial or automotive, it is very hard to be the technology leader. You may be the quality and reliability leader, but it's hard to be the leader [with regard to] the latest technologies, because that investment requires high volume and faster time to money.

Automotive markets are traditionally slower time to money. The fact that Micron has both pieces is a very strategic advantage for us because we can extend our leading-edge technology investments from mobile or storage markets to bring to market industrial- or automotive-quality and feature specific products. Our typical competition in these markets tends to be smaller mixed companies that are either a one- or two-technologies or product line type of company; they are not as broad a technology leader the way Micron is and that's what gives us a very good competitive advantage.

NEW! Intel® E3800 Series Edge-Connect Architecture

Applications

- *Unmanned drone/robotics* (air, surface, underwater) for leak detection, security, agriculture, science research, and energy
- *Mobile computing* for IoT, payload/mission computers, intelligent controllers or datalogging in a variety of rugged environments.
- *Portable* healthcare instrumentation and equipment.
- *Man-Wearable* Computing...especially in rugged use scenarios.

At just 75mm x 75mm the ADLE3800SEC is ideal for rugged, or extended temperature use in a variety of industries including: military, rugged industrial, unmanned, energy, transportation, medical or security and surveillance. **Edge-Connect Architecture** provides easy expansion and helps reduce cabling, integration time and system size all while increasing quality and overall MTBF.

ADLE3800SEC: E3845 Quad, E3827 DC

Front-Side I/O

- 2x 10/100/1000 LAN
- 1x USB 2.0
- 1x USB 3.0
- 1x DisplayPort

Mini Embedded PC: ADLEPC-1500

- Dimensions: 40mm x 87mm x 87mm
- I/O: 2x 10/100/1000 LAN; 1x USB2.0; 1x USB3.0
- 1x DisplayPort
- Options: 1xM.2 KeyB socket for PCIex1 or SATA SSD modules

Embedded Solutions
SMARTER BY DESIGN™

ADL Embedded Solutions
855-727-4200
sales@adl-usa.com

Designing with Intel® Embedded Processors?

Embedded Intel® Solutions delivers in-depth product, technology and design information to engineers and embedded developers who design with Intel® Embedded processors

Visit
www.embeddedintel.com

Subscribe Today at
www.embeddedintel.com
Free!

The Premier Conferences Devoted to Implementing IoT Technology in Embedded Systems

INTERNET OF THINGS DEVELOPERS CONFERENCE

The Big Event: April 26-27, 2017 | Santa Clara Convention Center

IoT Device Security Summit | September 14, 2017 | Santa Clara, CA

Connected Devices & Gateways Summit | November 30, 2017 | Santa Clara, CA

Register Now @

www.iot-devcon.com

www.congatec.us
6262 Ferris Square | San Diego CA 92121
Phone: 858-457-2600| sales-us@congatec.com

intel IoT Solutions Alliance

conga-B7XD

First COM Express Type 7 full compliant module.

Server-On-Module

- Intel® Xeon® D CPUs with up to 16 cores & 24 MB cache
- 32x PCI Express lanes, 2x 10 GBit Ethernet
- Smallest server board ever - 125 x 95 mm²
- Optimized for demanding real time applications
- Ideal for rugged micro servers for industrial environment

We simplify the use of embedded technology.