

# RTCA

Backplane Systems Remain  
the Workhorse of Industry

**16**

Automotive Systems Move  
to Software Standards


**32**

Multi-Level SSDs Move into  
Mobile and Rugged

**36**

THE MAGAZINE OF RECORD FOR THE EMBEDDED COMPUTING INDUSTRY

VOL 15 / NO 10 / OCT 2014

A hooded figure, possibly representing a hacker or a digital entity, is shown from the chest up, wearing a hooded garment. The figure is set against a background of blue binary code (0s and 1s) that appears to be floating or falling. At the bottom of the image, a portion of a globe is visible, showing the Earth's surface. The overall color scheme is dark blue and black with red text.

## Providing Security for the Wireless World



An RTC Group Publication



# Critical Recording in Any Arena When You Can't Afford to Miss a Beat!



**FREE**  
Talon SystemFlow  
Simulator  
Download Now!

Introducing Pentek's expanded line of Talon<sup>®</sup> COTS, rugged, portable and lab-based recorders. Built to capture wideband SIGINT, radar and communication signals right out-of-the-box:

- Analog RF/IF, 10 GbE, LVDS, sFPDP solutions
- Real-time sustained recording to 4 GB/sec
- Recording and playback operation
- Analog signal bandwidths to 1.6 GHz
- Shock and vibration resistant Solid State Drives
- GPS time and position stamping
- Hot-swappable storage to Windows<sup>®</sup> NTFS RAIDs
- Remote operation & multi-system synchronization
- SystemFlow<sup>®</sup> API & GUI with Signal Analyzer
- Complete documentation & lifetime support

Pentek's rugged turn-key recorders are built and tested for fast, reliable and secure operation in your environment.

Call 201-818-5900 or go to [www.pentek.com/go/rctalon](http://www.pentek.com/go/rctalon) for your FREE *High-Speed Recording Systems Handbook* and *Talon Recording Systems Catalog*.



**PENTEK**  
Setting the Standard for Digital Signal Processing



**28**

## Helping to Overcome Internet of Things Security Challenges with Wireless Infrastructure

### DEPARTMENTS

- 06 EDITORIAL**  
Security: Put What Efforts We Can toward the Problems that Count
- 08 INDUSTRY INSIDER**  
Latest Developments in the Embedded Marketplace
- 10 SMALL FORM FACTOR FORUM**  
Outside the Box
- 40 PRODUCTS & TECHNOLOGY**  
Newest Embedded Technology Used by Industry Leaders



**32**

## The New AUTOSAR Standard Is Reshaping the Automotive Landscape

### EDITORS REPORT

INTEL DEVELOPERS FORUM

- 12 Intel Debuts New Chips, Tools and Technologies to Link the IoT to the User**  
by Tom Williams, Editor-in-Chief

### TECHNOLOGY IN CONTEXT

BACKPLANE-BASED SYSTEMS

- 16 COM Express Benefits Extend Beyond Carrier Boards**  
by Greg Harrison and Earle Foster, Sealevel Systems, Inc.
- 20 Mobile Surveillance Systems: Leveraging the Traditional for the Design of the Future**  
by Lauren Wright, General Micro Systems

### TECHNOLOGY CONNECTED

SECURITY IN THE WIRELESS WORLD

- 24 Is Open Source Wireless Connectivity Worth the Security Risk?**  
by Dave Hughes, HCC
- 28 Helping to Overcome Internet of Things Security Challenges with Wireless Infrastructure**  
by Robert Day, Lynx Software Technologies

### TECHNOLOGY IN SYSTEMS

AUTOMOTIVE SYSTEMS: SMART, CONNECTED AND SAFE


- 32 The New AUTOSAR Standard Is Reshaping the Automotive Landscape**  
by Andrew Patterson, Mentor Graphics

### INDUSTRY WATCH

SOLID STATE MEMORY ADVANCES

- 36 Multi-Level Cell SSDs perform in HPEC Rugged Environment**  
by Steve Gudknecht and Ken Grob, Elma Electronic





**congatec**  
the rhythm of embedded computing

## Bridge the gap between ARM and x86 with Qseven Computer-on-Modules

One carrierboard can be equipped with Freescale® ARM, Intel® Atom™ or AMD® G-Series processor-based Qseven Computer-on-Modules.

conga-QMX6



ARM Quad Core

conga-QA3



Intel® Atom™

conga-QG



AMD® G-Series SOC

Additional details at: [www.congatec.us](http://www.congatec.us)

**congatec, Inc.**

6262 Ferris Square | San Diego | CA 92121 USA | Phone 1-858-457-2600 | sales-us@congatec.com



**MSC Embedded Inc.**  
Tel. +1 650 616 4068  
info@mscembedded.com  
www.mscembedded.com



## Qseven™ - MSC Q7-IMX6

Compatible Modules from Single-Core to Quad-Core

The MSC Q7-IMX6 with ARM Cortex™-A9 CPU is a compatible module with economic single-core CPU, strong dual-core processor or a powerful quad-core CPU with up to 1.2 GHz, and provides a very high-performance graphics.

- Freescale i.MX6 Quad-, Dual- or Single-Core ARM Cortex-A9 up to 1.2 GHz
- up to 4 GB DDR3 SDRAM
- up to 64 GB Flash
- GbE, PCIe x1, SATA-II, USB
- Triple independent display support
- HDMI/DVI + LVDS up to 1920x1200
- Dual-channel LVDS also usable as 2x LVDS up to 1280x720
- OpenGL® ES 1.1/2.0, OpenVG™ 1.1, OpenCL™ 1.1 EP
- UART, Audio, CAN, SPI, I2C
- Industrial temperature range

V-7\_2013-W0E1-6535

## RTC MAGAZINE



### PUBLISHER

#### President

John Reardon, johnr@rtcgroup.com

#### Vice President

Aaron Foellmi, aaronf@rtcgroup.com

### EDITORIAL

#### Editor-in-Chief

Tom Williams, tomw@rtcgroup.com

#### Senior Editor

Clarence Peckham, clarencep@rtcgroup.com

#### Contributing Editors

Colin McCracken and Paul Rosenfeld

#### Copy Editor

Rochelle Cohn

### ART/PRODUCTION

#### Art Director

Jim Bell, jimb@rtcgroup.com

#### Graphic Designer

Michael Farina, michael@rtcgroup.com

### ADVERTISING/WEB ADVERTISING

#### Western Regional Sales Manager

Mike Duran, michael@rtcgroup.com  
(949) 226-2024

#### Midwest, Canada, EMEA and Asia Sales Manager

Mark Dunaway, markd@rtcgroup.com  
(949) 226-2023

#### Eastern Regional Advertising Manager

Jasmine Formanek, jasminef@rtcgroup.com  
(949) 226-2004

### BILLING

#### Vice President of Finance

Cindy Muir, cmuir@rtcgroup.com  
(949) 226-2021

### TO CONTACT RTC MAGAZINE:

#### Home Office

The RTC Group, 905 Calle Amanecer, Suite 250,  
San Clemente, CA 92673  
Phone: (949) 226-2000  
Fax: (949) 226-2050  
Web: www.rtcgroup.com

#### Editorial Office

Tom Williams, Editor-in-Chief  
1669 Nelson Road, No. 2,  
Scotts Valley, CA 95066  
Phone: (831) 335-1509  
tomw@rtcgroup.com

#### Published by The RTC Group

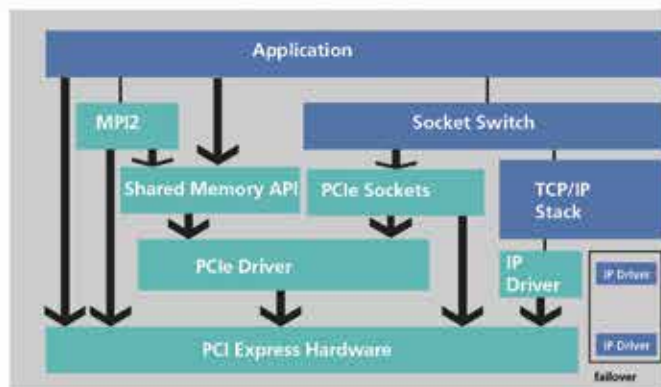
Copyright 2014, The RTC Group. Printed in the United States. All rights reserved. All related graphics are trademarks of The RTC Group. All other brand and product names are the property of their holders.

# Need Software for PCI Express®



## PCI Express® Network

PCI Express® Networks are the right choice for high speed embedded applications. Dolphin software supports IDT and PLX non-transparent bridging chipsets. If you are building a custom board or system products, Dolphin PCI Express® software can dramatically reduce your development time and improve system performance.



### PCI Express® Software

Dolphin PCI Express software is a complete software stack that supports Windows, Linux, and now VxWorks. This stack includes software for Peer to Peer connections, sockets, reflective memory connections, and TCP/IP support.



# Security: Put What Efforts We Can toward the Problems that Count

BY TOM WILLIAMS, EDITOR-IN-CHIEF

For those of you who are regular readers of this space, you may be aware that when it comes to questions of security, I am at best skeptical. And you don't want to hear me when I'm at less than best. One primary reason is that as a matter of just plain fact, there is no such thing as absolute security. It is always relative to the skill, patience and motivation of the attacker. That said, there are certainly levels of security that can be attained at a certain level of price and effort. While these do have value, there should be no illusion that they will always protect the precious data jewels they were designed to guard.

One example is the recent introduction of Apple's iOS8 operating systems, which now supposedly puts all the user data that is on the device behind an "unbreakable" pass code. What that really seems to mean is that Apple has designed the system so that there is no way for even Apple to bypass the passcode (oh heck, let's just call it what it is—the password) to get at the user data. So the password protects access to the device. Well, that's pretty much the same as before. The difference is that Apple itself is technically prevented from bypassing that password on orders of a search warrant because it has been made technically impossible. Impossible, that is, without the password.

Remember, still, we are talking about a given device running iOS8. If you have put all your email, photos and plans for attacking government installations up on iCloud, however, they are every bit as accessible as before and Apple would have to

comply with warrants and subpoenas as before. The upshot is that the normal user is pretty much as protected against the nongovernmental hacker as before. I can't remember the last Chinese intelligence officer who showed up at Apple's door with a search warrant.

What users of iOS8 are really protected against is Apple itself and Apple is well motivated to stress that point, however subtly because as far as we know, Apple has never scooped up user data to sell or use for targeting advertising. In contrast to their arch-rival, Google. So in some sense, the whole hoopla about privacy seems to be aimed at making Apple the good guys and taking a swipe at Google. This new protection is really not going to significantly affect many people. If somehow you can get the password, you can still get in.

This, of course, has not stopped the orgy of hoopla that broke out on the mass media about "game changers" and "concern in the intelligence community." All the while that same mass media is bringing almost daily reports about huge data breaches in places like Target, Home Depot, Goodwill and more and about a massive credit card hack by Russia. The ones that don't make it into the mainstream media are really the scariest—rogue cell phone towers, attacks on power grid infrastructure and more. The never-ending battle needs to be fought where it really counts.

When we give up the idea of absolute security or of universal security, we can focus on robust targeted security. This is especially needed in

the era of the Internet of Things and despite many concerns that may be expressed, there are truly enormous efforts and advances taking place. Security is not free and robust security requires both money and effort. The problem with securing the IoT is that there are so many different devices and levels of concentration that a uniform approach to security is well-nigh impossible.

Small devices with sensors, large machines with sensors, actuators and other sources of data feed into aggregation points and gateways that eventually lead up to the Cloud but through different paths and different levels of complexity. Selecting exactly where and how to apply the best security effort is difficult at best.

In addition, we are still searching for the best methods and technologies for implementing IoT security at different levels and at costs that are comfortable for those levels. There is certainly no lack of hardware and software products and services claiming to do just that, but there is insufficient experience to really be able to tell what will work best. As we can see from recent events, the experience we are getting is coming at quite a cost. And we have to realize that security is not a problem that can be solved; it is a problem that must be managed. It is a constant struggle that will only improve as we apply new methods from new information and learn to minimize the losses while enhancing the benefits. And in many cases, it won't be pretty.



## OSS Flash Storage Arrays

Visit One Stop Systems  
SC14 booth #754

### Direct Attached Flash Storage



**FSA 25** 1U chassis with 25TB Flash storage attaches to a single server.



**FSA 50** 2U chassis with 50TB Flash storage attaches to one or two servers.

Attach to each server at 128Gb/s  
PCIe x16 3.0 cable adapters and cables included



**FSA 200** 3U chassis with 200TB Flash storage attaches to one to four servers

### Flash Test Systems

Test up to 32 PCIe cards or nVME drives at a time  
Open canisters provide easy access, power cycling, and voltage margining of PCIe cards or nVME drives.



Test  
nVME SSD's

Test  
PCIe Flash boards

### The World PLC Market Faced Another Decline for 2013

After a downturn in 2012, the global PLC market declined again in 2013, with revenue falling by 2.1 percent annually. As the most mature market, Europe had the largest base for PLC sales in terms of revenue. However, it is very slowly dragging itself out of recession. With tight fiscal policies, tight credit conditions in several countries, excess industrial capacity and still relatively weak export demand, there are few signs of a strong upturn in the near future. In light of the lack of strong domestic market demand, the PLC market had encountered a small decline as a whole.

The U.S. market for PLCs is positive; the housing market continues to recover, consumer spending is rising, export markets are improving and the pace of capital spending is increasing. Because of that, the overall market for PLCs in the United States is growing strongly. Looking at end-user sectors, the fastest growing segments are the oil and gas and food and beverage industries.

Asia Pacific, which has been one of the fastest growing regions in the past few years, had faced a much slower growth rate than before. China's market is the most likely reason for that. China's leadership has singled out financial stability as its most important objective, with only moderate stimulus applied. Because of that, many investments have been delayed. Besides, China—as an export-focused market—had faced weakening demand from its leading trading partners, such as Europe. The PLC market had faced two-sided pressure from both domestic and foreign markets.

The Japanese market, however, had a good year. The Japanese government had pushed forward strong economic incentive plans since 2012 by applying fiscal stimulus, engaging in monetary easing and implementing structural reforms. But when turning into the U.S. dollar, the market showed a decline in growth because the currency had depreciated by more than 10 percent in terms of exchange from the Yen to the U.S. dollar.

## Researchers Pioneer Spray-On Solar Cells; Propose Diverting Lead Waste for Low-Cost Photovoltaics

In a discovery that could help cut the cost of solar electricity, a team of scientists at the University of Sheffield has fabricated perovskite solar cells using a spray-painting process. The researchers had used the spray-painting method previously to produce solar cells using organic semiconductors—but using perovskite is a major step forward, they asserted.

Efficient organometal halide perovskite-based photovoltaics were first demonstrated in 2012 and are now a promising new material for solar cells as they combine high efficiency with low materials costs. This class of material offers the potential to combine the high performance of mature solar cell technologies with the low embedded energy costs of production of organic photovoltaics, the researchers explained.

In a complementary development, a system proposed by MIT researchers recycles materials from discarded car batteries, which are a potential source of lead pollution, into new, long-lasting solar panels that provide emissions-free power by using lead to produce perovskite—specifically, organolead halide perovskite—a technology that has rapidly progressed from initial experiments to a point where its efficiency is nearly competitive with that of other types of solar cells.

Initial descriptions of the perovskite technology identified its use of lead, whose production from raw ores can produce toxic residues, as a drawback. But by using recycled lead from old car batteries, the manufacturing process can instead be used to divert toxic material from landfills and reuse it in photovoltaic panels that could go on producing power for decades. And because the perovskite photovoltaic material takes the form of a thin film just half a micrometer thick, the MIT teams' analysis showed that the lead from a single car battery could produce enough solar panels to provide power for 30 households.

The best certified efficiencies from organic solar cells are around 10%, where perovskite cells now have efficiencies of up to 19%, which is not so far behind that of silicon at 25%—the material that currently dominates the worldwide solar market.

## Global Market for Smart Machines to Reach \$15.2 Billion by 2019

A smart machine is a machine that can accomplish its designated task in the presence of uncertainty and variability in its environment. However, not all smart machines are physical devices such as industrial machines and autonomous vehicles. Indeed, the market also comprises intelligent agents, virtual reality assistants, expert systems and embedded software that make traditional devices “smart” in a very specialized way. Increasing R&D spending, technological advances and consumer demand will drive tremendous growth in this market for the foreseeable future, according to a report from BCC Research. *Smart Machines: Technologies and Global Markets* provides an in-depth analysis of the global market for smart machines. According to the report, this market was valued at \$5.3 billion in 2013 and is expected to reach \$6.2 billion by 2014. BCC Research projects the market to grow to \$15.2 billion by 2019, and register a five-year compound annual growth rate of 19.7% from 2014 to 2019.



# Diablo Technologies and Supermicro Team to Develop Next-Generation MCS Enabled Systems

Diablo Technologies has announced a strategic partnership with Supermicro to bring the industry's lowest latency, non-volatile memory solution to server systems. As part of the initiative, Supermicro customers will have access to the latest X9-series platforms optimized for Memory Channel Storage (MCS), through the SanDisk ULLtraDIMM SSD. Supermicro's extensive Green Computing portfolio (including Twin architecture, GPU compute, SuperStorage and Hyper-Speed hardware accelerated platforms) now features a broad selection of MCS-enabled solutions. In addition, Diablo and Supermicro will collaborate on next-generation server and storage architectures, targeting a wide range of mission-critical, enterprise workloads.

Supermicro's SuperServer and SuperStorage architectures deliver significantly advanced levels of integration between system memory and NAND flash. The company's Green Computing solutions provide flexibility to scale out MCS-based devices to customized levels of capacity, performance and acceleration. Applications such as virtualization, big data analytics, database and low-latency trading can now seamlessly benefit from the disruptive performance advantages provided by the MCS architecture. As part of the integration, MCS-enabled systems will be available to Supermicro customers worldwide with support for major operating systems, including Linux, Windows and VMware ESXi 5.1 and 5.5.

## Himax Technologies and Lumus Collaborate to Develop Next-Gen Smart Glasses

Himax Technologies and Lumus have announced a joint initiative to continue developing the next-generation of smart glasses that will set new technological standards in image quality and performance.

According to Zvi Lapidot, chief executive officer of Lumus, "Himax's superior LCOS technology, its availability for high volume production, and the company's forward looking technological applications, were critical in our selection of Himax as a strategic partner. Their microdisplay, specifically designed for smart glasses, combines smoothly with Lumus' transparent display, creating the ideal solution for true augmented reality and hands-free wearable computing."

Mr. Lapidot added, "While our ultra-thin, see-through optics enable natural looking wearable displays, Himax's unique LCOS technology provides the high level of brightness necessary for see-through augmented reality. Ultimately, our cooperation enables us to bring widely appealing solutions to help seamlessly and intuitively blend wearable technology into our daily lives."

Himax and Lumus have been successfully collaborating for several years in the field of combat aviation, producing market-leading helmet mounted displays. Leveraging their combat-proven solutions and manufacturing capabilities, the two companies are now collaborating to make wearable display mainstream consumer products.

## Cmosis Acquires AWAIBA, Maker of Industrial and Medical Image Sensors

Cmosis has acquired all outstanding shares of the Swiss AWAIBA Group. AWAIBA develops and markets innovative line-scan CMOS imagers for industrial web inspection, sub-mm CMOS camera modules for endoscopy, and onboard automotive cameras.

"The acquisition of AWAIBA is a natural fit for Cmosis. It strengthens our presence in existing markets and expands our activities in adjacent segments with complementary image sensor products. It also strengthens our relationship with our existing customer base," stated Luc De Mey, Cmosis Chairman and CEO.

"The acquisition became possible after TA Associates stepped in as a strong financial partner of Cmosis earlier this year. This enabled horizontal expansion and growth acceleration," Mr. De Mey continued. "AWAIBA is a profitable and well-established brand with an attractive and strongly growing client base. We are very excited to welcome their experienced team, having an impressive track record of innovation and deep understanding of customer needs. This permits us to even better serve our customers as an independent and pure-play supplier of CMOS image sensors."

A report published by IC Insights in July of this year predicts image sensor markets exceeding \$13.2 billion by 2017. Cmosis and AWAIBA are focusing on the high-growth segments such as machine vision, production cameras, traffic enforcement, medical devices and prosumer products. According to IC Insights, these segments will account for about \$3.3 billion in 2017. CMOS technology continues to gain market share over charge-coupled device (CCD) technology.



# OUTSIDE THE BOX

WRITTEN BY COLIN MCCRACKIN

Start with the external connectivity. End with the peripherals and network interfaces, full circle. In between, select one or more processors and chipsets to crunch the data before passing back through the OS to some internal storage or out across cables or through radios.

As embedded engineers, we have to admit that we often fall for processor vendors' marketing pitches. We've got to have the latest and greatest high-speed machine. We want to build in enough processing headroom for future features and software upgrades without going back later to change the hardware. We want to brag about the cutting-edge design on our resume. And of course our sales team tells us that customers like the sizzle of certain processor brand names inside the box. Clearly the hype would have us approaching our projects in the reverse direction: inside-out.

Your mission, should you choose to accept it, is to design an optimal solution while filtering out the noise and spin of other people who have an agenda. To do this, resist the temptation to choose the latest whiz-bang processor first. Tell the pointy haired boss who has attended too many seminars to take a vacation instead. If he tells you to "think outside the box," tell him that you already are, both figuratively and literally. That ought to get him or her off your back, for long enough to

finally understand your point at least. Embedded system designs must start and end with the external operating environment. After all, it's the I/O, dummy. (No, we haven't forgotten that it's also the software, dummy.)

This year is another banner year for processor launches. In the embedded x86 world, performance seems to grow faster than can be consumed by the operating systems and applications. Detractors of Moore's Law watch yet again as aggressive die shrinks and bleeding edge transistors and dielectrics provide a glut of processing power and bus bandwidth. The single core 2 GHz processors from several generations back with only generation 1 of PCI Express are now easily outperformed by dual and quad core processors with gen 2 lanes from the affordable embedded ultra-mobile-based roadmap. Larger cache sizes and faster RAM interfaces for the cache misses complete the picture. New low power microserver processors are applauded. If the processing can be distributed or offloaded from the main CPU, lower power consumption would mean reducing overall size and weight of the device we are designing. Depending on the co-processing alternatives, the cost may come down too.

After analyzing I/O requirements and scaling down the main processor, then examine the impressive array

of standard form factor industrial computer boards and processor modules. If the Mini-ITX shoe fits, wear it. If the I/O circuitry is available in standard slot cards and the shock and vibration requirements are modest, an inexpensive industrial ITX solution can be cobbled together. If too bulky or flimsy, a custom carrier board can implement the I/O with the exact layout and connectorization desired. This design path leads to a carrier that can be reused for generations with a simple compatible CPU module swap. The overall process leads to good results, and there is simply too much at stake—including your reputation as a designer—to jump the gun.

The amount of money that can be saved by downsizing the processor next time is substantial. It could reduce the size, weight and cost of the thermal solution as well. Before gulping down the processor vendor's excess transistor Kool-Aid, take a close look at what's just enough to keep the I/O happy. Start outside the system, trace the bandwidth inside the box, through the CPU and back out again. Your buyer and accounting team will become your new best friends. And your boss should come around too, once the Kool-Aid buzz wears off.



# THE POWER INSIDE TOMORROW'S TECHNOLOGY



PCIe-IDIO-24

## PCIe-IDIO-24 PCI EXPRESS DIGITAL I/O CARDS

- 24 optically isolated, non-polarized digital inputs
- Polarity insensitive AC/DC inputs accept up to 31 VDC or AC RMS
- Software configurable filters on inputs for electrically noisy environments
- Can detect input state change and assert interrupt
- 24 optically isolated fully protected FET high-side switch outputs
- Outputs capable of switching from 5-34 VDC at 2A



IGEPv5

## IGEPv5 ARM-BASED SINGLE BOARD COMPUTER

- OMAP5432 ARM Cortex-A15 Dual Core up to 1.7GHz
- More than 3.5 Dhrystone MIPS / MHz
- Up to 170 million polygons/second graphics
- HD video capable (1080p encode/decode)
- Up to 4GB DDR3 RAM and 8GB eMMC Flash (with full version)
- 2D GC320 Core and 2D/3D Dual Core SGX544 Graphic Accelerators



## INDUSTRIAL GRADE FLASH STORAGE

- mSATA, sSATA, 2.5" SSD, SD, SDHC, CFC, CFast, UFD
- Advanced universal wear leveling and block management
- High endurance
- Intensive TDBI for high reliability
- In COM and IND Temp grades
- Industry best Data Loss on Power Fail
- Exceptional resistance to shock and vibration



Centronics to USB  
Print Converter



V2 Slimline  
USB Floppy Emulator

## LEGACY PRINTING SOLUTIONS & FLOPPY EMULATORS

USB Print Converters: *Centronics Print Converter (RS-232 also available)*

- 1x Centronics female (printer input)
- 2x USB 1.1, 1x LAN 10/100 Ethernet

USB Floppy Emulators: *Replace Floppy Disk with USB*

- DOS/Windows compatible
- 1.44MB floppy disk size
- Read/Write directly to the USB pen drive



THE EMBEDDED PRODUCTS SOURCE  
1.800.548.2319 ➤ [www.wdlsystems.com](http://www.wdlsystems.com)  
[www.wdlsystems.com/mag/RTC/oct14.html](http://www.wdlsystems.com/mag/RTC/oct14.html)

**WDL**  
SYSTEMS  
[sales@wdlsystems.com](mailto:sales@wdlsystems.com)

# Intel Debuts New Chips, Tools and Technologies to Link the IoT to the User

The recent Intel Developers Forum gave engineers and marketers alike a glimpse of some new possibilities for mobile, low-power devices for industry and consumer that will dwell in a world of greater connectivity and fewer wires.

by Tom Williams, Editor-in-Chief

The Intel Developers Forum (IDF) that was just held in San Francisco was, of course, a big venue for Intel to show off its latest goodies and advocate its vision of the current direction of embedded technology. Having said that, it was also a genuine opportunity to get a sense of where development and markets are headed in the coming year. Companies like Intel (and there are admittedly few of those) not only set trends; they also respond to them.

These trends include, naturally, the Internet of Things and all that entails in terms of low-cost, low-power silicon, connectivity, servers, small modules and an increasing attention to end user and consumer needs. These include enhanced smartphones, tablets and PCs and wearable computing—all of which are well-known buzz words.

However, the first day of the conference offered an interesting contrast and helped to highlight Intel's role in the industry in which it is such a huge player. On September 9, Apple held a conference in Cupertino, CA, at which it introduced two major products—the iPhone 6 and the Apple watch. Of course, the press was all over this nationwide while the Intel conference did get some coverage in the local San Francisco paper. The major difference is that Apple designs not only the underlying silicon and software functionality of its end-user products, it also meticulously crafts the actual end products down to the last detail of form and style.

That same day, Intel announced an agreement with Fossil Group, which is known for its watches, but also for other fashion accessories, to “identify, support and develop emerging trends in the wearable technology space.” We can assume that the first products to emerge from this alliance will be watches but perhaps other wearable accessories as well. The point is that Intel supplies the underlying technology and for the consumer space, partners like Fossil are responsible for the pizzazz. Also, it does not exclude other partners, who may also

be competitors. And that, not surprisingly, was why the focus of the conference was on developers and partners.

## Processor Platforms

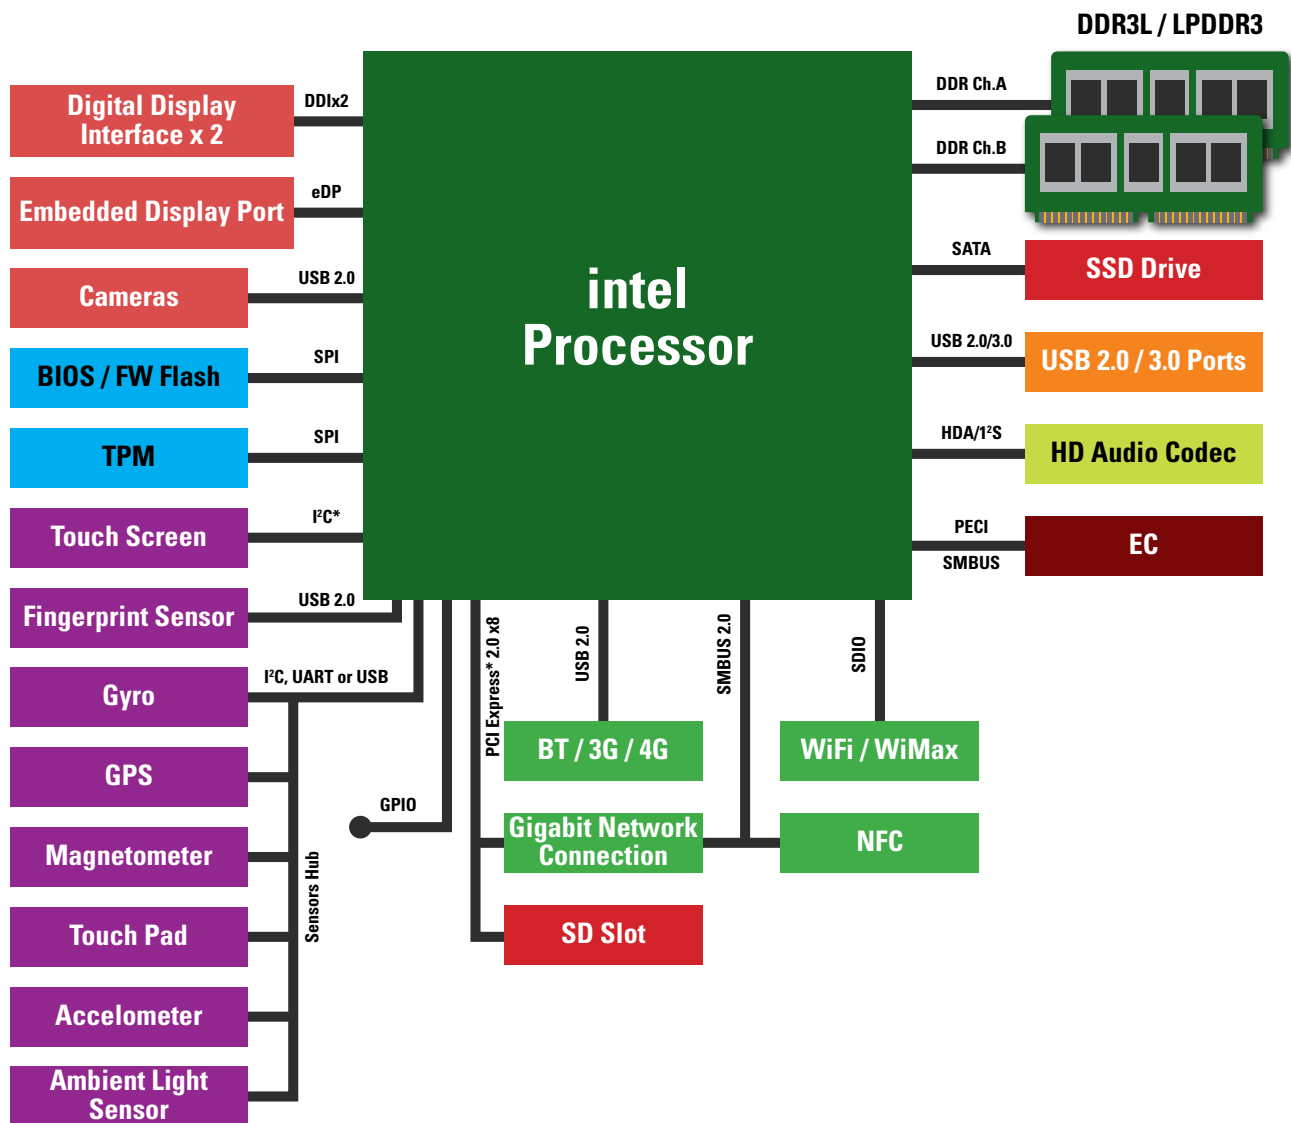
Among this year's stars is the Core M Processor family, which currently consists of three 64-bit multicore processors, the 5Y10, 5Y10A and the 5Y70 build on Intel's 14nm technology. A separate die inside the processor package called the platform controller hub (PCH) supplies the processor family I/O. This includes interfaces for sensors such as gyros, accelerometers, GPS and more that are increasingly found in mobile devices (Figure 1).

While mobile tablets and laptops along with 2 in 1 tablet/laptop devices are the initial targets of the Core M, there is little doubt that with its power consumption coming somewhere between the Atom and earlier Core families and a 50% boost in speed and an approximate 40% increase in graphics performance, it will be finding its way into a vast array of embedded applications, both industrial and consumer.

In an additional aid to developers, Intel also announced its Edison development platform, which is based on a 22nm technology dual core Atom SoC, formerly Silvermont, on a board that is just 35.5 mm x 25 mm. The Edison—with a recommended customer price of \$50—supplies interfaces in the form of 12 GPIO, I2C, UART, SPI, USB 2.0, 6 analog inputs and a clock output and is aimed at the development of small IoT and wearable computing devices. The Edison will initially support development with Arduino and C/C++ followed by Node.JS, Python and later by visual programming tools.

While Intel has long had the reputation of producing processors in first instance for the PC and laptop markets and then also targeting them for the embedded arena, its SoCs that include Silvermont and Baytail with their many on-chip peripherals, internal buses and rich I/O really do appear to be





**Figure 1**  
The Intel Core M processor includes a die in the package that supports a platform controller hub (PCH) that provides rich I/O.

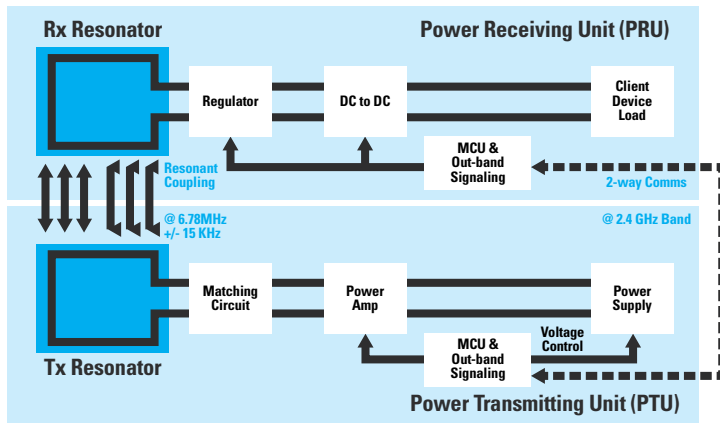
aimed squarely at the embedded developer. And they certainly are rapidly being adopted by that segment of the design community.

### The PC's Descendants and the IoT

The evolution of the PC does appear to continue to be an important focus for Intel. Even as tablets, smartphones and the Internet of Things explode, there is still a need for a user's central access point to applications and personal data. Interestingly, there is also a need to pay attention to the design

and configuration of servers that will be needed to house the enormous amounts of data generated by the Internet of Things and the various connected devices from wearable nodes to phones, autonomous control and monitoring systems and PCs that increasingly can take the form of small, thin, powerful notebooks or 2 in 1 devices, or "portable all-in-ones," that can work as a touch tablet or a PC with keyboard.

To enable partners to develop such a next generation, Intel is bringing out its Core M processors mentioned above to be followed by the next fifth generation of 14nm Core i5 and i7 vPro processors and the Core i3, i5 and i7 devices in early 2015. The scope of these introductions including SoCs like Silvermont and its supporting Edison board pretty clearly points to targeting the full range of applications from small, wearable devices to portable and mobile machines on up to the



**Figure 2**  
The Rezenze charging technology uses resonant coupling at 6.78 MHz to charge client devices. Bluetooth Low Energy matches the power transmitting unit (PTU) to the characteristics of the power receiving unit (PRU).

full world of gaming to be addressed by 4 GHz devices and the Core i7-5960X 8-core “extreme” processor. And they’re going to want to do it all without wires. And also without wires, they are going to want sophisticated graphics and video including 3D with facial recognition capabilities—all this they aim to put into mass market consumer devices.

**Getting Rid of Wires**

Among the oncoming wireless technologies is WiDi, a wireless display technology developed by Intel that allows streaming of display and video from a portable device to a larger display or an HDTV. Intel Pro WiDi also has a security feature that puts a privacy screen up on both the presenter’s PC and a conference room screen, for example, so that a Intel display can be shared with a trusted group. WiDi is currently supported by Intel’s fourth generation Core vPro processors and beyond.

Intel is also supporting the WiGig multi-gigabit wireless technology for such things as wireless docking and high-speed data transfer. WiGig was developed by the WiGig Alliance, which has now been subsumed by the Wi-Fi Alliance. The technology is capable of transfer rates up to 7 Gbit/s, although it typically cannot transmit through walls at that speed. It is not intended as a replacement for Wi-Fi but as a supplemental technology that can be useful at short ranges.

And then there is wireless charging of devices, which is based on a technology called Rezenze and supported by the Alliance for Wireless Power (A4WP). The user experience of Rezenze technology involves using a metal plate below almost any surface to enable the charging of any Rezenze-enabled device. Such surfaces can be set up anywhere such as in desks

and tables as well as public places like retail stores, airports or even office lobbies.

The wireless power transfer (WPT) system transfers power from a single power transmitter unit (PTU) to up to eight power receiver units (PRU’s.) The power transmission frequency is 6.78 MHz. The system also supports a Bluetooth Low Energy (BLE) link that is intended for control of power levels, identification of valid loads and protection of non-compliant devices. The PTU comprises three main units, a resonator and matching unit, a power conversion unit, and a signaling and control unit. The PRU also has three functional units like the PTU (Figure 2).

The control and communication protocol for the WPT network is designed to signal PRU characteristics to the PTU as well as to provide feedback to enable efficiency optimization, over-voltage protection, under-voltage avoidance, and rogue object detection. The WPT network is a star topology with the PTU as the master and PRUs as slaves. Thus the end user need only set the target device on the surface and the protocols link it up for automatic charging.

Intel is being fairly specific about the classes of device it sees its technologies targeting, but history has shown that the embedded industry traditionally takes advantage of technologies that had been initially introduced for the PC and mass market (such as USB, PCIe, SATA, etc.) and adopts them for all manner of specialized and unique embedded devices and applications. It will be no different with such technologies as the newer class of processors, the connectivity technologies and more. Now that the Internet of Things is increasingly connecting the consumer with background industrial systems and processes, we can just imagine what the new generation of devices and technologies will bring.

**Intel**  
**Santa Clara, CA.**  
**(408) 765-8080**  
**www.intel.com**



# RUGGED C4ISR SYSTEMS SINCE 1979

## AIRBORNE SYSTEMS

### “GOLDEN-EYES” SB1002-MD

**Rugged, Dual, Fully Isolated Systems with Removable Drives**

- Fully supports Multi-Domain, NSA-secure architecture
- Two fully independent systems (Red/Black) **each with:**
- Up to 2.4GHz Intel® Quad Core™ i7 Haswell with 6MB of L2 Cache
- Up to 32GB of 1600MHz DDR3 memory with ECC
- Up to 1TB of sealed removable nDrive SSD drives and 1TB fixed
- Ultra-low SWaP, 5.4” x 6.5” x 2.75” @ 6 lbs and as low as 40W total



## PORTABLE SYSTEMS

### “MARLIN” SG502-LP

**Fully Sealed, Rugged, Ultra-Low-Power System**

- 2.0 GHz Intel® Quad Core Atom™ with 2MB of L2 cache
- Up to 4GB of 1333MHz DDR3 SDRAM with ECC
- Up to 1TB fixed high speed SSD
- One Gigabit Ethernet port with Power over Ethernet (optional)
- Flexible high performance I/O configuration options
- Ultra-low SWaP, 6.0” x 3.75” x 1.0” @ 1.5 lbs and under 15W



5.56mm M855A1 EPR shown for scale\*



## VISION SYSTEMS

### “RUGGEDVIEW”

**Rugged, Touchscreen Smart Displays with Removable Drive(s)**

- Ultra-rugged and lightweight, less than 2 inches thick!
- Up to 2.4GHz Intel® Quad Core™ i7 Haswell with 6MB of L2 Cache
- Up to 32GB of 1600MHz DDR3 memory with ECC
- Available in 10”, 15”, 19” in 4:3 format
- Available in 12”, 17”, 24” in 16:9 format
- Available in 32”, 55”, 65” in 4K 16:9 format



## VETRONIC SYSTEMS

### “TARANTULA” SO302-4in1

**Rugged, Secure Virtual Machine, 18 port Switch, RAID and APU**

- Up to 2.4GHz Intel® 10 Core Xeon® Ivy Bridge-EP with Hyper-Threading
- Up to 128GB of DDR3 RAM with ECC up to 1600 MT/s
- Up to 6 Secure Virtual Machines and 18-port managed layer 2/3 switch
- Up to 16TB removable canister storage with hardware RAID controller
- Auxillary Power Supply (APU) for orderly system shutdown
- Ultra-low SWaP, 11.75” x 7.7.5” x 4.5” @ 18 lbs



**GENERAL MICRO SYSTEMS, INC.**

PROUDLY DESIGNED & MANUFACTURED IN THE U.S.A.

[www.gms4sbc.com](http://www.gms4sbc.com)

8358 Maple Place, Rancho Cucamonga, CA 91730 • (909) 980-4863 • (800) 307-4863



# COM Express Benefits Extend Beyond Carrier Boards


Many backplane standards are available that offer varying electrical and mechanical features, but custom systems can also benefit from a backplane architecture. COM Express, a popular implementation of Computer on Module (COM), offers a powerful processing platform with several bus connectivity options perfect for creating versatile backplane systems.

by Greg Harrison and Earle Foster, Sealevel Systems, Inc.

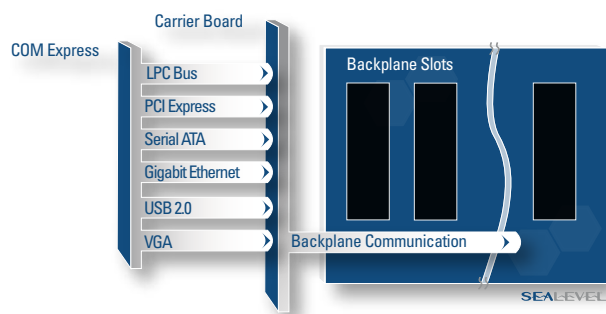


**Figure 1**

cExpress-HL Compact COM Express Type 6 Module from Adlink offers a 4th generation Intel Core i7/i5/i3 processor. COM Express modules provide core computing functionality and connect to a carrier board that provides connectors and application specific I/O.







**Figure 2**  
Block diagram shows a COM Express module connected to a carrier board and backplane. COM Express offers a variety of interfaces suitable for backplane communications.



COM Express, a popular solution for creating custom systems, offers designers numerous advantages including fast time-to-market, scalable processing, flexible mechanical options and long product life cycle. The COM Express processor module delivers the high-speed functionality germane to all computers such as video, disk interfaces, memory, etc. The module typically connects to an application-specific carrier board that provides the exact configuration of I/O and connectors for a specific purpose. COM Express architecture seems attractive for OEM products with a set functionality (Figure 1), but what about applications that require configuring the I/O set to match various installation locations?

For example, an OEM instrument designed for collecting environmental data at remote sites will usually need to be configured with the I/O types and count specific to each location. Designing a custom carrier board that includes all possible I/O configurations is likely to be size and cost prohibitive. These applications scream for a backplane system to provide the required flexibility. Backplanes also afford room for adding functionality when new application requirements arise.

Developers shouldn't overlook COM Express and the benefits it provides as an option for a backplane solution. The COM Express standard defines multiple connectivity options that are brought out via the module-to-board interconnect, including USB, Ethernet and PCI Express. Alternatively, RS-485 and other familiar serial interfaces can be easily generated from USB or PCIe buses on the carrier board and routed to I/O positions on the backplane.

Depending on the desired system mechanical configuration, the I/O slots and supporting hardware can reside directly on the COM Express carrier board or can be supported on a separate

backplane PCB. Each interface choice presents advantages and trade-offs. Designers can choose the best backplane communication configuration by considering the attributes of each interface, the number of slots required, and the throughput/response time required for I/O interface (Figure 2).

## Backplane Communication Options

PCI Express, widely implemented as the expansion bus on modern PCs, is the most obvious choice for interfacing backplane I/O with COM Express modules. PCI Express is organized in serial "lanes," each consisting of a low-voltage differential TX/RX pair. COM Express provides up to six PCIe lanes for general purpose use (there is a separate group for graphics). Each single lane has a data rate of 2.5 Gbit/s, and multiple lanes can be grouped together to increase bandwidth, although this is not generally required for accessing general purpose I/O. The high data rate combined with physical-layer flow control capabilities makes PCIe suitable for reliable communications even in high-bandwidth applications.

Still, six lanes can be a limiting factor in backplane design, especially if any lanes are used for creating system functionality, such as serial ports on the carrier board. To expand beyond the six lane limit, a PCI Express packet switch can be implemented on the carrier board or backplane to create the desired number of backplane slots. However, more PCIe lanes results in a bandwidth penalty, so a hybrid approach that implements some number of dedicated PCI Express slots for high-speed backplane I/O while using one of the other serial interface options for the remaining general purpose slots may be more effective.

PCI Express, like all high-speed signals, requires care in layout in order to maintain signal integrity across the backplane. For best performance, signals must be routed as matched-length differential pairs, use controlled impedance, and stay within the electrical length requirements of the standard.

USB can also be a good choice for communicating with I/O cards over a backplane. The COM Express specification requires at least four USB ports, but many modules include the maximum of eight. Some of these ports likely need to be available to the system as general purpose ports, but by implementing USB hubs on the carrier board or backplane, the desired number of ports can be assigned to backplane slots.

USB supports automatic enumeration for detection of USB connected I/O cards and the installation of necessary software drivers. Although USB uses differential signaling, routing USB signals over a backplane requires care since the standard encodes single-ended state information. USB 2.0 offers three speeds for communications: 480 Mbit/s high speed, 12 Mbit/s full speed and 1.5 Mbit/s low speed.

Although the data rates are relatively high, the USB specification requires the host to poll USB devices. This polling architecture, along with other factors including processor performance, operating system, application software and amount of I/O affect latency, makes USB unsuitable for real-time or near real-time





**Figure 3**

Sealevel Systems' Relio R3 rackmount industrial computer uses a COM Express to backplane architecture. Relio R3s offer 19 expansion slots for application-specific I/O.

requirements. However, for general purpose I/O, USB is a good choice for backplane connectivity.

COM Express modules do not typically include serial port functionality, but implementing RS-485 communications over a backplane is straightforward using the COM Express module's USB or PCIe bus and simple circuitry on the carrier board or backplane. Long a popular standard for industrial communications I/O, RS-485 uses differential signaling that offers noise immunity suitable for electrically noisy environments and is well suited to routing over a backplane.

RS-485 offers relatively low-speed communications, up to 921.6 Kbit/s, and the standard defines an address for each I/O location that simplifies software protocol development. As an alternative to creating a custom software driver to handle communications, industry standard Modbus RTU offers a well-defined, documented option supported by a variety of third-party software packages.

Ethernet is widely used for backplane communications in a wide range of products designed for military and commercial applications. COM Express modules supply a minimum of one Ethernet port, and the specification recommends 10/100/1000BaseT Gigabit for the port. This port can be connected to the backplane slots by adding Ethernet switch circuitry to the carrier board or backplane.

Ethernet implements flow control primarily in software, resulting in possible data loss when there is a large amount of traffic. In that case, adding a managed switch circuit can often

improve performance adequately for applications that do not require real-time response. For applications that require deterministic timing, industrial Ethernet protocols like EtherCAT work best.

## Rackmount System Example

Sealevel's Relio R3 industrial computer uses a COM Express engine to power a 3U 19" rackmount computer with a total of 19 I/O slots. One PCIe slot is included for high-speed I/O such as video processing, while the other 18 slots connect via RS-485 to Sealevel SeaRAQ I/O boards. As shown in Figure 3, the COM Express processor mounts to a carrier board that brings out all the standard features of the COM Express module including DisplayPort video, Ethernet and USB channels. RS-485 is generated on the carrier board, and a transition board holds the PCIe connector and routes the RS-485 signal to the vertically mounted backplane PCB. This architecture allows configuring the Relio with a choice of Intel i7, i3, or Atom processor simply by changing the COM Express module. As technology evolves and faster processors are available, the system can be upgraded with minimal engineering effort by changing the COM Express module.

For applications that require real-time response, smart I/O backplane cards can offload functionality from the host processor. This approach can eliminate backplane throughput and data integrity concerns for any of the COM Express connectivity choices. Advances in Field Programmable Gate Array (FPGA) technology make designing smart I/O cards easier than ever (Figure 4).



**Figure 4**

Sealevel SeaRAQ I/O boards connect to the COM Express processor via RS-485. For high-speed applications, FPGAs can be used on I/O boards to reduce throughput requirements.



An FPGA is a good choice for interfacing a PCIe lane from the COM Express module to I/O circuitry on a backplane card. By creating a PCIe endpoint and custom I/O interface logic in the FPGA, the application running on the COM Express module can quickly and easily read and write data using a memory-mapped addressing format. The result is an elegant software interface with fast system I/O response time.

For real-time requirements, it is now possible to run an operating system and execute an application program written in a high-level language directly in an FPGA on an I/O card. A number of CPU cores are available for use in FPGAs. For example, the NIOS II "soft processor" from Altera provides flexibility, tight integration with the FPGA logic, and the ability to easily create custom I/O peripherals. The NIOS II core resides in the FPGA fabric and provides a full 32-bit processing engine capable of running an operating system and allowing application software to be developed in a high level language.

The COM Express backplane architecture provides the freedom to exactly match the application I/O and mechanical requirements while providing an easy upgrade path for the core processing functions that are most likely to change, thereby extending the useful life cycle of the system. OEMs will benefit most from the investment required in the initial design through the product's configurability, scalability and long life cycle.

Additionally, end users will benefit from the ability to upgrade existing hardware with minimal costs compared to a full system replacement.

**Sealevel Systems**  
**Liberty SC.**  
**(864) 843-4343**  
**www.sealevel.com**



# The CUBE™ expansion enclosures

ONE STOP SYSTEMS

Choose from a variety of options:

ExpressCard, PCIe, or Thunderbolt connectivity package	1, 2, 3, 5, or 8 slots	Full-length (13.25"), mid-length (9.5"), or short card (7.5")	Half-height or full-height cards	36W, 180W, 400W, 550W or 1100W power supply
--	------------------------	---	----------------------------------	---

**Flexible and Versatile:** Supports any combination of Flash drives, video, film editing, GPU's, and other PCIe I/O cards.

**ORDER TODAY!**

The CUBE, The mCUBE, and The nanoCUBE are trademarks of One Stop Systems, Inc. Maxexpansion.com and the Maxexpansion.com logo are trademarks of One Stop Systems, Inc. Thunderbolt and the Thunderbolt logo are trademarks of the Intel Corporation in the U.S. and other countries.



**Figure 1**

The nature of surveillance requires a system that can ensure the reliability of complex components working together.

# Mobile Surveillance Systems: Leveraging the Traditional for the Design of the Future

The demand for ruggedness, small size, secure mass storage, sensors, displays, low power and high connectivity in today's mobile surveillance systems continues to grow. Only by using the latest compact, powerful components and subsystems can these growing demands be met.

BY LAUREN WRIGHT, GENERAL MICRO SYSTEMS





**Figure 2**  
The small stature of the comprehensive system (as compared to a full sized 5U rackmount) saves space without sacrificing function or performance.

At the core of successful surveillance is the ability to collect data without detection. This can be problematic once those under observation become privy to the current, most innovative technologies used. Thus, as this current era of innovation evolves, the difficulties involved in collecting, interpreting and processing ever-increasing amounts of data, sometimes in intense mobile environments, place extreme pressure on mobile platform designers and engineers. Current trends and future projections show that surveillance efforts continue to require greater consideration for systems technology in regard to size, weight and power (SWaP), as well as user experience and high-security capabilities. That said, as organizations like the Border Patrol and other national security agencies gear up for the inevitabilities of the future, the pressure to obtain surveillance technologies that overcome existing and forecasted roadblocks is now of an especially time-sensitive importance.

A unique difficulty presented during surveillance efforts, particularly in mobile situations, involves the consistent exertion to guarantee seamless interoperability between multiple systems. As several independent technologies are linked and utilized for detection and collection initiatives, the user must ensure that they all work in concert in order to avoid a mission-critical catastrophe. After all, multiple points of failure exist between numerous recording cameras, night vision functions, data storage, displays, sensors and network communications. If an integral component (e.g., a network line) fails, the user could be left with several valueless heaps of heavy metal until the system is repaired, wasting time and aborting opportunities for critical data collection.


These challenges are exacerbated in mobile environments where vehicle operation and safety are as important as the mission itself. Traditionally, cameras are attached to the exterior

of the vehicle, absorbing data and showing the driver what to expect from the outside environment via inboard monitors and sometimes requiring a night vision device (NVD) for functionality. Because the driver of the vehicle must be able to navigate his expedition using a networked vision system, there is even greater pressure on both the reliability of the vehicle and computational accuracy. Until recently, even the most advanced technologies only offered vision systems with unsuitably high latencies, inflicting depth-related motion sickness and a misaligned reality between what the driver sees and the actual location of the vehicle.

Part of what helps create a low latency vision system is the speed at which the computer's network operates. As all cameras must be connected to one another and able to communicate easily with a mainframe database, high-speed networking and data transfer capabilities are incredibly important for surveillance applications (Figure 1).

For greatest efficiency, the system should have the ability to capture video at a rate of 30 frames per second using up to 16 HD-PTZ, analog data cameras. The system then converts the data to a GigE Vision format with lossless compression, streams the data over Gigabit Ethernet from each camera to a router/switch, and then sends it to the server via 10 Gigabit Ethernet, at speeds that ensure no information is lost. The captured video is then displayed, in broadcast mode, on any of the interior smart display monitors via Gigabit Ethernet or 10 Gigabit Ethernet. The video is stored on a storage subsystem, all in real time without losing any frames or enduring latencies of more than one frame. That would be a very long process for a computer that isn't "up to speed," especially considering that the network is also tasked to shift between additional sensors and other communication and computational systems throughout the vehicle.

Whether collecting data for analysis, judicial processes, or other critical applications, ensured data security rests with the recording device being used to store that information. As quickly as cameras and other sensors are able to collect information, a device must be used to safely accumulate and store the uncompressed data for future analysis. This requires incredibly accurate recording



**Figure 3**

The many elements of the SO302-4in1 Tarantula work in concert for a fully efficient surveillance system.

capabilities. Should the device incur a hiccup during the recording process, the potential for data corruption increases. This threat simultaneously increases the significance of reliable high-security recording devices, making the choice of what system to use a difficult and calculated one. Storage devices configured with security mechanisms, such as AES encryption, secure erase and write protect, are required during surveillance operations because they establish protection in any environment.

Preferably, every aspect of the operation should have some sort

of embedded precautionary failsafe, as unanticipated circumstances are more than likely to arise. In the event of an unexpected power outage, for example, it is necessary for the vehicle to include a component that allows the internal systems to undergo a self-sustaining, orderly shutdown. This is particularly vital during surveillance efforts, as an uncontrolled shutdown can result in severe loss of acquired data. This specialized device can come in the form of an auxiliary power unit or uninterruptible power supply.

Ranging from commercial SUVs to military ground transport, a clear challenge for system designers is presented through the very limited amount of space available in surveillance vehicles. Until recently, the standard computing technologies used in these vehicles were VPX and older backplane platforms, which are characteristically bulky and require a considerable amount of energy.

To conceptualize the breadth of space and power consumed by these platforms from a commercial perspective, imagine an F150 with a 1U server and processor from Dell, a 1U managed switch from Cisco, a 2U storage with NAS capabilities, and a 1U auxiliary power unit. That's a full 5U rackmount application requiring well over 5,000 watts and 12-15 times the necessary volume when compared with stand-alone systems. The consequences of this type of arrangement result in heavy, large, hot systems that utilize an exceptional amount of space and necessitate an effective method for heat removal.

Ideally, a fully integrated, independent system that includes secure storage capabilities, intelligent I/O, monitor support for a vision application, and ultra-fast networking to tie it all together is the best solution for a highly functional, mobile platform. In order to reduce the threat of multiple points of failure between interconnects, adhere to SWaP constraints, and ensure data security, mobile platform designers have been tasked with creating this all-inclusive solution (Figure 2).

### Many Functions Little Space

Upon conceptualizing a stand-alone system for these mission requirements, it was realized that an increasing amount of additional components needed to be included. Cameras, communication, radios, displays, possible weaponry, positioning systems, sensors and hydraulic systems are all necessary for proper surveillance techniques, but fitting them all nicely into a commercial sized vehicle is a daunting task. However, by adding multiple virtualized workstations that enable sophisticated processing and multi-channel intelligent I/O, the various independent technologies on the vehicle can be supported while saving space and power. A large selection of highly flexible I/O that can sustain multiple internal monitors in order to match the performance of the computational communication equipment is also essential.

An example of one such integrated system is General Micro Systems' S0302-4in1 (Tarantula). The backbone of this product is an Intel Xeon Ivy Bridge-EP CPU with 10 cores (2.4 GHz each) driving six independent virtual machines and controlling up to 18 Gigabit Ethernet ports and a second 10 Gigabit Ethernet port. It also contains up to eight 2 Tbyte SATA SSD drives (16 Tbyte total) in one canister with RAID capabilities and an internal APU in another canister. The APU is comprised of an array of super capacitors that provide power per MIL-STD-704 blackout requirements.

One of the more severe environments where such a system can be employed is the U.S. Army's MRAP Night Vision Program. The fortified MRAP vehicles used in the program are built to rove rural, mountainous and dangerous environments with one intent being stealthy surveillance and detection. As urgency prevails, MRAPs must often use night vision devices to navigate hazardous areas in the dark. This requires several cameras attached to the outside of the vehicle that send signals to the displays inside. Accurate estimations of ground topography are critical in these situations, as the amount of delay in communication between the

cameras and the displays could mean life or death.

For these reasons the Army has chosen to use the Tarantula, a system that supports GigE Vision protocol to provide a low-latency of ½ to 1 frame for its vision application. In other words, from the time it takes for the camera to see something to the time the data is processed internally, there is less than 1 frame of video delay. This is also made possible through the networking communication within the system, as it fully supports managed layer II and layer III functions, such as VLAN and QoS processing, enabling differentiated services delivery and security through intelligent frame processing and egress frame manipulation. The MRAP vehicles include 17" and 12" internal monitors, also provided by General Micro Systems, which display the external cameras' video playback. These touchscreen smart displays use gigabit or 10 Gigabit Ethernet for extreme speed and data processing. They also include bezel keys that are used to determine which camera to view at any given moment, and a night vision imaging system (NVIS) for use in low or no light situations. Industry standard GigE Vision also allows for pan, tilt and zoom capabilities, providing the commander of the vehicle and passengers the ability to see any potential hazards and the environmental orientation of the area (Figure 3).


The Army also required that the system include six individual virtualized workstations with separate I/O in order to control real-time video, defensive counter measures and other critical operations. Each of the six I/O sites is fully independent, connected to the host CPU via PCI Express lanes only, meaning that all I/O of one workstation is separate from the I/O of another, and they are all fully monitored for security through Trusted Platform Module (TPM) and Trusted Execution Technology (TXT).

Moreover, it was vital that the system also be equipped with embedded security measures, such as tamper-proof protection, which recognizes unfamiliar access of software and BIOS boot and locks the system, only allowing restart with controlled reauthorization. Another key requisite allows an authorized user to "zero-ize" the system, placing all data and programs at zero for information fortification. These mechanisms are embedded within the system as precautionary elements to aid in the preservation of any sensitive information obtained during surveillance applications or otherwise.

The Army's program highlights one of the more acute applications that these stand-alone systems are being employed for. However, small systems like these will be key to all future mobile platforms requiring sophisticated processing, vision and communication. Today's intense reliance on electronic surveillance and data collection will only increase, and providing smaller, more powerful systems that utilize less power will continue to test the talents of contemporary system developers.

**General Micro Systems**  
**Rancho Cucamonga, CA.**  
**(909) 980-4863**  
**[www.gms4sbc.com](http://www.gms4sbc.com)**

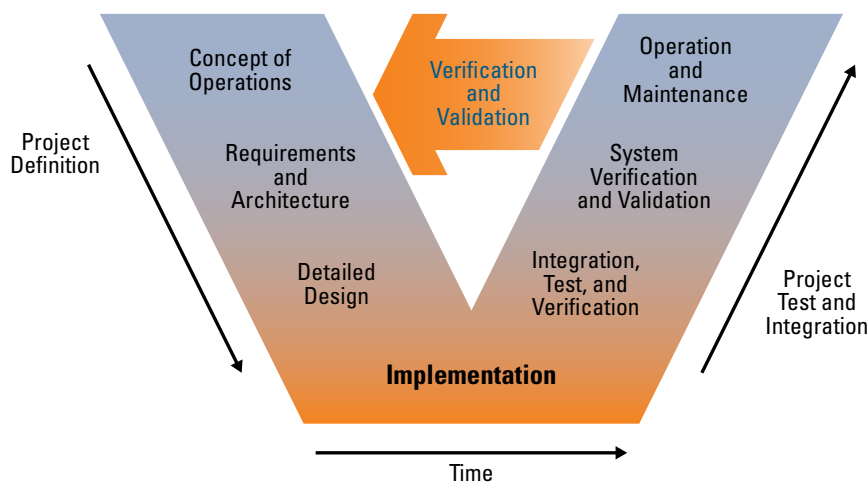




## Is Open Source Wireless Connectivity **Worth the Security Risk?**

The Heartbleed security breach, based on OpenSSL, raises the spectre of attacks across a range of wirelessly connected embedded devices. Rigorous software development processes are critical for protecting wirelessly connected devices in the Internet of Things.

by Dave Hughes, HCC



**Figure 1**  
V model of the systems engineering process from: "Systems Engineering Process II" by Osborne, Brummond et al.

Open Secure Sockets Layer (OpenSSL) is widely used to provide network security in many different kinds of computing systems, including wirelessly connected embedded systems in the emerging Internet of Things. OpenSSL is also the open source security library that allowed the widely publicized security breach called Heartbleed. While there are advantages to open source libraries such as OpenSSL, there are clearly risks as well, many of which stem from the development process itself. The main process used for development of OpenSSL is simple. First a programmer develops code, then a reviewer checks the code, and finally the code is released.

This method of development is the way most software in the world is developed. If you look behind the scenes of OpenSSL development, there are usually four programmers, only one of whom is full time. This leads to a fairly obvious question—why do huge companies, often with access to significant engineering resources, trust their customers' data and their own reputation to such a small team; especially to a team outside of their control, which may potentially expose the company to unquantifiable quality and security risks? "Because it has always been done that way" would seem to be an insufficient response in the light of the chaos caused by Heartbleed.

In retrospect, Heartbleed seems to be more of a warning tremor than a full earthquake. It showed the potential scope and depth of harm, but the consequences of this particular fault were relatively mild. Continuing to follow the same path, however, will undoubtedly lead to similar problems, and the ubiquity of the software is in itself a weakness, which can be exploited by those who choose to do harm.

### Better Software Development Methods Needed

If the methods of development used by OpenSSL were demonstrably the state-of-the-art in robust software development, then there would not be much to debate. However security problems such as Heartbleed, Apple's "goto fail" and GnuTLS have been caused by defects in software, not necessarily in the protocols or design. Across various industries there are well-established methods for developing high-quality software. The aerospace, industrial, medical and transport industries use software processes based on the "V" model development defined by IEC 61508, and the data shows that not only does it reduce defects significantly, but in many cases it also reduces the cost of software management over its lifecycle.

How would use of such methods have helped in the OpenSSL Heartbleed bug case? Let's

A TQMP2020 module with a Freescale QorIQ can save you design time and money



TQ embedded modules:

- Are the smallest in the industry, without compromising quality and reliability
- Bring out all the processor signals to the Tyco connectors
- Can reduce development time by as much as 12 months

The TQMP2020 module comes with a Freescale QorIQ™ Power Architecture® MCU and supports Linux and QNX operating systems.

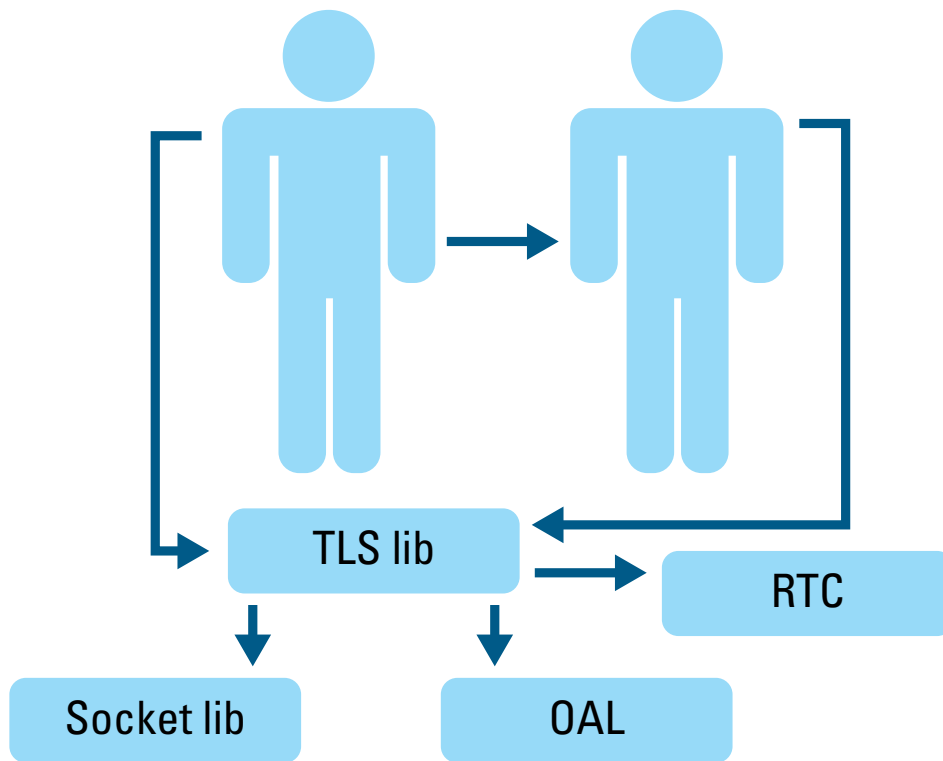
The full-function STKP2020 Starter Kit is an easy and inexpensive platform to test and evaluate the TQMP2020 module.



Technology in Quality

ConvergencePromotions.com/TQ-USA

TQ-USA is the brand for a module product line represented in N. America by Convergence Promotions, LLC



**Figure 2**

Developers should choose development processes that reflect how much they value the security of customer data.

look at some specific development approaches that can help address security specifically.

### “V” Model Development

In the Heartbleed situation, the information available states that the software failed to check the scope of a protocol variable and then processed it blindly. Standard V model development would include unit testing and boundary case analysis/testing that would have instantly alerted developers to the issue (Figure 1). There are other elements of the process that would also have picked up these kinds of issues. For example, a decent static analysis tool would have picked up Apple’s recent issue with their TLS software.

It would be impractical from either a cost or resource point of view to propose that full V model development be used for all software, and it is not the intention of this article to state that open-source methodologies are “bad.” Open source software is open, not just in the source code but also in the processes used to develop this software. It is no secret in the industry what processes could be used to achieve a low software defect level. However, no open-source software today goes to these lengths, and in the area of security the question is—is this approach

good enough? Indeed, can it ever be good enough?

### Verification of Software Components

When a company wants to use any piece of equipment in a highly sensitive application area, you would expect the manufacturer of that equipment to verify that all components used reach the required level of quality. It is unclear how this occurs in companies managing large amounts of potentially sensitive customer data. This always happens in a manufacturing process where they check the supplier history, the strength of components, ISO9001 compliance, etc., but strangely not for security.

There are deeper issues to consider. If it were possible to create a perfect TLS implementation, would that mean the system was secure? More secure maybe, but if a defect bug was sitting elsewhere in the target system (e.g., in the TCPIP stack), then it could be possible to expose memory. It is much less likely that it would yield sensitive data, but still possible. Eventually we conclude it is necessary to ensure every part of a sensitive system is designed to a verifiable standard.

The only practical solution is to carefully partition what belongs in a critical part of a system and what does not. For instance, bringing the whole of Linux (used in many of the systems that use OpenSSL) to this standard is clearly unrealistic. There is a risk that someone could make a mistake in a Linux update, or BSP update, in unrelated code and leave systems vulnerable. This would leave very little possibility for companies basing their products on this system to protect themselves.



## The Problem with “Free” Software and Security

If we accept that mistakes will always be made and systems will tend to become more complex, then continuing as things are now will probably result in further problems. Commercial devaluation of software does not help this process. The idea that software can be created and obtained for free is a bizarre concept for commercial companies to believe in. It also appears to focus only on the initial capital cost of software and not the ongoing maintenance costs. If the lifetime cost of development and maintenance of “free” software was truly accounted for, it would probably raise some corporate eyebrows.

It could also be quite difficult for any company involved in a “Toyota style” legal case where the consequences of software errors were much worse than compromised data. Imagine a defect, caused by a mistake by a hobby programmer in Australia and reviewed by a programmer working in his spare time in Argentina, which resulted in injury or loss of life.

Again this is not an attack on open source—they are open and transparent. Blind usage of any software without a proper assessment of context and risk is the problem. Developers should choose development processes that reflect how much they value the security of customer data (Figure 2).

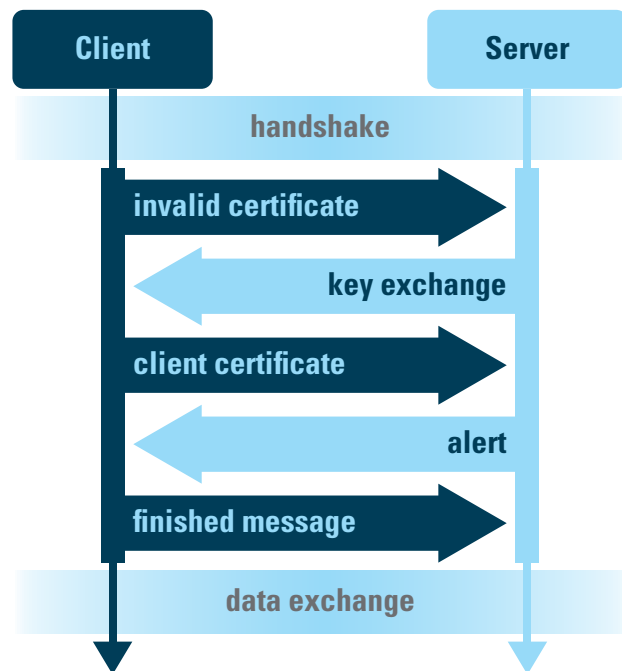
The argument that software is open and therefore everyone will fix everything is clearly not sustainable anymore—the Heartbleed bug existed for two years before someone realized the problem. This would not be acceptable in any safety-critical or secure environment. There are several different issues.

First, the only way it was possible to exploit the Heartbleed bug was by challenging a system that used OpenSSL. High security systems (weapons systems, nuclear power stations, etc.) publish as little information as possible about the system internals to make the attacker’s starting point as difficult as possible. There are practical problems with transport layer security (TLS) in this respect since the point is secure interoperability. Therefore the communication protocols used must be in the public domain. But OpenSSL is so widely used that, if an issue is discovered, it is relatively easy to find a victim. Concealing the details of an implementation reduces the likelihood that an attack can be effectively mounted.

Second, attacks on TLS-related algorithms in recent years have revolved around back-door methods, such as changes in power consumption or response times, rather than hacking the algorithms directly. These attacks are normally only possible with a direct knowledge of the specific algorithm used. In the TLS case again there are limitations because the algorithm to be used must be publicly negotiated. However, the specific realization does not need to be public knowledge.

## Moving to Secure Embedded Software Components

The commercial market for standard software components has been damaged by free software from many sources. How this



**Figure 3**

Formal development methods are well understood and will reduce the likelihood of security issues caused by software defects.

affects professional companies who need good quality code and support is not obvious. It seems that developers lose the benefits of scale that using specialist providers brings. HCC, an embedded software vendor, has always focused on high quality, reliable components, such as failsafe file systems, but we are working on components developed to standards of verifiability. Ultimately many of these will achieve certification under the IEC 61508 SIL3. We strongly believe that key components of embedded software should be developed once and reused in many environments. Providing these components with the necessary life-cycle support and documentation can make this level of quality more affordable across the industry.

The security of devices has become a critical issue for both device manufacturers and consumers. Wireless embedded devices have specific security issues based on their applications, though a large part of making them secure requires a rigorous approach to the development of software for them. As in similarly sensitive fields such as aerospace and medical, a formal approach to development will significantly reduce the probability of a major incident with a product (Figure 3).

**HCC Embedded USA**  
**New York, NY.**  
**(212) 734-1345**  
**[www.hcc-embedded.com](http://www.hcc-embedded.com)**



# Helping to Overcome Internet of Things Security Challenges with Wireless Infrastructure

Some key embedded security technologies can be used both in the IoT endpoints and sensors, as well as in the IoT infrastructure to provide a defense in depth against tomorrow's cyber threats. It is important that these measures be incorporated as the network is established and not as afterthoughts.

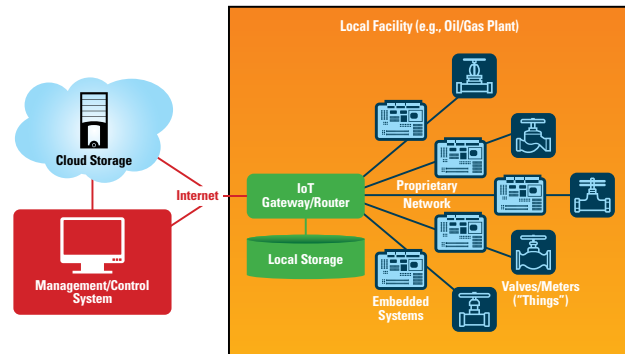
by Robert Day, Lynx Software Technologies

The Internet of Things (IoT) is driving toward ubiquitous connected embedded devices that can be potentially accessed from anywhere in the world. Wireless connectivity is also helping to fuel that drive by removing many of the traditional barriers (cost, complexity, physical installation) that were borne by physical network connectivity. Although wireless connectivity is helping accelerate our connected world, it is also helping to open our new IoT world to the increased likelihood of cyber-attacks, especially around our critical infrastructure. A defense in depth strategy must be employed as we put this IoT infrastructure in place, as traditional network and endpoint security is not adequately containing today's cyber threats.

Before looking at the security challenges, it's a good idea to review a typical IoT network topology, which will help explain where the potentially vulnerable parts of an IoT connected system lie. We will use an example topology that represents the industrial automation world, often called SCADA (supervisory, control and data acquisition), which will use a computer infrastructure around energy generation, including smart grid management and also oil and gas refineries and distribution systems. The relatively recent cyber-attacks on some foreign infrastructure, including the infamous Stuxnet virus that managed to infect and control an Iranian Nuclear facility, show just how vulnerable even the most secure, fortified and remote systems can be.

The embedded devices or "things" are typically linked to a physical item and have the primary function of communicating with that physical item. That communication might be read-only (i.e. monitoring) or read/write (monitor and control), and the data that is communicated can be read or initiated by either humans or machines (M2M). These physical systems can include smart meters and electricity control breakers for power control, or valves and flow meters for oil and gas. Either way, the quality and security of the data is paramount to the reliable function of the system. These embedded devices are networked (to receive or provide data), but are typically not linked directly to the Internet, as they are normally connected to a proprietary network that is usually local to the facility where the "things" are physically located. This could be an electricity substation, or a whole oil and gas plant, depending on the scale of the system. In the days before IoT, this was a relatively secure network, as access could only be obtained by being physically present at the site, and so could be contained using physical security measures (security guards, barbed wire, etc.). As the cost and convenience of wireless networking has spread to these local networks, the physical security measures may not be quite as effective, since hackers could reside outside of the physical site and still gain wireless access to the network. So effective protection is required on both the embedded devices and the local network, but this is still relatively low on the security risk spectrum, compared to where this data goes next (Figure 1).

For our ubiquitous IoT world, the data from these embedded devices now needs to be aggregated and fed to the people or machines that will use this data. This could include manage-



**Figure 1**  
An example Industrial IoT network topology.

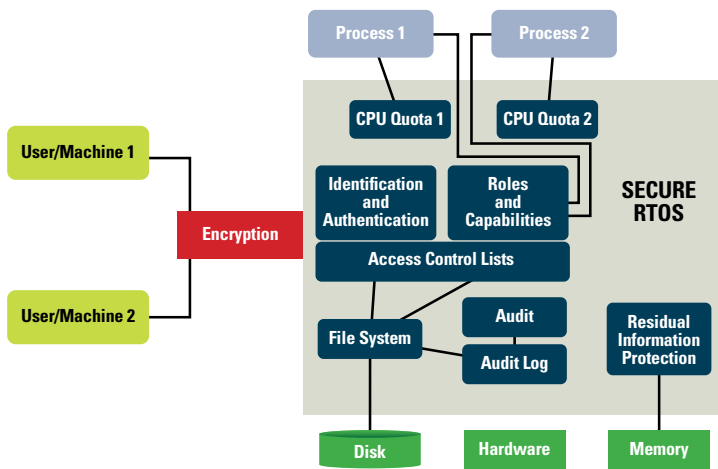


ment and billing for the electric grid, and plant logistics, yield management, safety and control for oil and gas plants. Assuming that these consumers of the data are not physically located on site, this data will need to speed its way to another location, typically using the same Internet that everyone else in the world is connected to. Quite how much data leaves the site really depends on the application of the data, and also on how much local intelligence, aggregation and storage is available.

The connection to the Internet is generally achieved using a local gateway or router in much the same way that a local home router gives both Internet access and local Wi-Fi networking. These IoT gateways have to connect to all the local embedded devices (on their proprietary network) and then to the outside world using the Internet. Although they are often physically located on site, they are very open to the outside world through the Internet connection. This is a potential security nightmare, as anyone who gains access to the router, has the keys to the physical kingdom without having to be physically present. Depending on how much intelligence is in the router, a lot of data analysis and aggregation could be done here, even though its primary function is still connecting and communicating data between two networks.

After the router has decided what data needs to be sent on, the data is generally encrypted and sent to the next destination. This is often a large data storage facility, often housed in the Cloud, where data analytics can be used to provide meaningful information on the data, such as billing or usage data, maintenance data, or yield and performance information. Alternatively, the data is sent to some management and control systems where actions are taken either by humans or machines to control the embedded and physical devices at the site. So where are the vulnerabilities? They are typically not in the data in transit, as it is encrypted, but are





**Figure 2**

A secure RTOS used to secure the embedded system connected to the “things.”

often at the final destination of the data, when it is decrypted and hence vulnerable to either theft or compromise.

Looking again at this somewhat simplified IoT network, it is easy to see that there are some attack points that really need to have good security. The embedded device itself is vulnerable either to local wireless attacks, or potentially to attack by a compromised router. The router itself is a huge potential attack point, as it is connected to both the local network and the Internet, and without some secure separation between the two, this could be a very easy place for remote cyber-attacks. Finally, the Cloud storage or remote management systems are also potential attack points, with a much larger potential payoff as they hold all the data (from all the embedded devices at all the sites), and will often have control and override functions at both a site and/or device level, plus they are all connected to the attack-prone Internet.

## Defense in Depth

So, a defense in depth strategy needs to be implemented, to help protect all the vulnerable parts of the network from all types and methods of cyber-attack. Luckily, technology is available that if used when the network is being designed (rather than as an afterthought), could dramatically reduce the chance or effect of an attack. Much of this technology has evolved to meet the security needs of the Department of Defense (DoD), which has been operating secure remote networks for decades, and where a compromise in any part of the network could be fatal to national security and hence not an option.

Most of the world’s security functionality has been implemented as add-ons on top of existing infrastructure, or as patches to help seal security gaps in the infrastructure. As an example, think of protection that one needs for a regular desktop or laptop PC; antivirus software, firewalls, OS security patches,

not to mention all the application security additions, and the vast amounts of network security appliances that surround the network infrastructure trying to thwart cyber-attacks, usually by looking for attacks that are similar to previous ones. In the IoT, rather like the DoD, one attack could be fatal, and therefore, preventative security needs to be built into all parts of the infrastructure.

Similar to the PC world, operating systems are often the key attack point as they are typically the highest privileged software in any given system, and if compromised offer keys to the control and data kingdoms. So there needs to be a better way to protect these operating systems than the traditional anti-virus or OS patch mechanisms that are normally used. Operating systems are used throughout all of the IoT infrastructure topology described above, and we need to look at security solutions for each of these parts separately. We will focus on securing the parts of the infrastructure that are typically considered embedded systems, the devices and the routers, as the Cloud and management systems have security issues that are generally serviced by IT security products and vendors.

First, the “things.” These things are connected embedded systems, often not using large operating systems, but using hard real-time OSs (RTOSs) that are helping to support the networking function and the data extraction or control function of the physical item they are connected to. These RTOSs have been traditionally more secure than desktop OSs, often because of their proprietary interfaces, but also due to the fact that they haven’t been as connected to the outside world as they are becoming with the IoT. So, adding wireless network connectivity makes these “things” a lot more vulnerable, and their proprietary interfaces will not stop a determined attacker who has gained entry via the wireless network. However, if an RTOS is used that has built-in security functionality, especially one that was designed to meet the exacting security needs of DoD tactical systems, then it could offer enough security protection to stop the most determined attacker.

Examples of operating systems with built-in security include Security Enhanced Linux (SELinux) and the LynxOS RTOS. Both of these operating systems introduce a number of key security concepts that help the OS protect against malicious attacks regardless of how they enter the system. These concepts can include discretionary access control for file system objects, fine-grained user access control using roles and capabilities, identification and authentication control of users, device and system quotas to help thwart DDoS attacks, trusted path mechanisms for guaranteed communication links, and residual information protection to stop attacks by reusing or viewing used memory (Figure 2). An RTOS with these built-in security features is the best protection for the embedded wireless device, as it still offers the real-time characteristics, supports the required network functionality, typically has a smaller footprint than a GPOS like Linux, and now offers advanced protection.

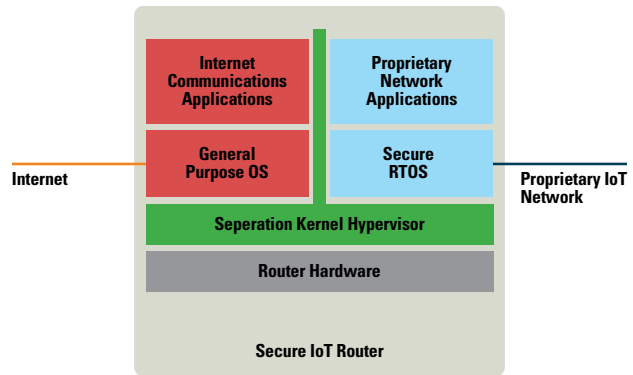
Secondly, the router. This is a bit more challenging to protect

as we are now dealing with a much more complex system that has to support multiple networks (including wireless), needs to connect securely to the untrusted Internet, and at the same time is passing and processing a large amount of data. Routers will often use an OS that is more complex than the traditional RTOS but still gives a security issue as it has a lot of software functionality and a large attack surface being controlled by a single privileged software entity. A secure OS as suggested for the embedded systems is a good step in the right direction, but due to the complexity, multiple networks, and its link to the Internet, this really needs to be stepped up a notch, and leads us more toward what is known as a multi-domain system.

In the DoD, secure OSs have been used for multi-domain systems linked to different networks at different levels of security classification, but the prevailing thoughts and technology for true domain separation call for something known as a separation kernel. This is at a higher level of privilege than the OS (i.e., it sits between the OS and the hardware), and its primary function (as the name suggests) is to separate the resources in the system, such that an attack in one domain cannot reach or compromise the other domain. In order to still offer the functionality required from an OS, the separation kernel also contains virtualization functionality that allows the “guest” OS (or OSs) to reside above it in separated secure virtual domains. This separation kernel approach gives some very interesting benefits when designing these highly intricate cornerstones of the IoT.

Firstly, security. The small separation kernel is the only software item at the highest privilege level, and if designed properly it will not contain untrusted elements such as device drivers or software stacks, as they can now reside in the lower privilege guest OSs. This substantially reduces the “attack surface” of the highest privileged software. Any attacks made on the guest operating systems will be contained in their own secure domain, without compromising the rest of the system, which essentially stops the attack from spreading and likely reaching its intended target. This is key to protecting the proprietary network and the “things,” as the most likely attack point is through the Internet, and that operating system is not connected directly to the proprietary network, so the keys to the IoT kingdom are safely locked away in their own domain (Figure 3).

Secondly, flexibility, suitability and performance. By having multiple guest operating systems in their own secure domains, we can now choose which OS best suits which domain. Before virtualization, a single OS had to control all the tasks in the router, and that generally meant adding general purpose functionality to an RTOS, or sacrificing real-time performance by using a GPOS. Now a GPOS can be used to connect to the Internet side of the router, and an RTOS (maybe the secure RTOS as described above) can be connected to the proprietary side. The two sides can only communicate with each other by using the secure internal networking channels provided by the separation kernel, which can be carefully moderated, controlled and in some instances made to be unidirectional.



**Figure 3**  
An IoT router/gateway securely protected by a separation kernel hypervisor.

In summary, the infrastructure that enables the Internet of Things is very vulnerable to cyber-attacks, especially as it embraces modern communication technologies, such as wireless networking and the Internet. And the more critical the infrastructure, the larger the threat. Energy companies specifically need to be very vigilant in securing their infrastructure, as a widespread attack here could render cities, states and even the country helpless. However, embedded software technology such as secure OSs and separation kernels, which have been helping to secure military infrastructure, are now available to help protect the IoT as it becomes more ubiquitous.

**Lynx Software Technologies**  
San Jose, CA.  
(408) 979-3900  
www.lynx.com



# The New AUTOSAR Standard Is Reshaping the Automotive Landscape

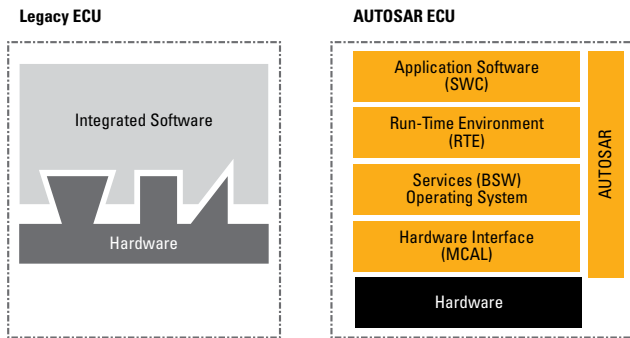
AUTOSAR provides a predefined standard approach for vehicle network and ECU design that is finding its way into every automotive OEM and Tier 1 organization. Here are some of the expected business benefits of an adoption strategy and a look at some of the basic terminology and design methods.

by Andrew Patterson, Mentor Graphics

Since its formation in 2003, the AUTomotive Open System ARchitecture (AUTOSAR) alliance has been changing the way vehicle networks and electronic control units (ECUs) are designed. AUTOSAR offers an industry standard approach for OEM manufacturers and their Tier 1 suppliers to design and develop ECUs that are at the heart of modern vehicles. The standard helps reduce the opportunity for human error in the design process and offers suppliers and manufacturers a well-defined, machine-readable data format for exchanging design information.

The AUTOSAR alliance has among its members automobile OEMs and a supporting ecosystem of





**Figure 1**  
Separating application software from hardware.



component and service providers. The goal of the alliance is to create and establish global open standards for automotive Electronics & Electronics (E/E) architectures. The standard assists at the vehicle architecture level, allowing OEM network designers to design and manage the complex interaction of vehicle functions, and also at the supplier level where details of individual ECUs interfaces need to be specified prior to manufacture.

### Why Make the Switch to AUTOSAR?

A modern luxury vehicle can contain up to 100 ECUs ranging in function from simple sensor interfaces to complex infotainment and telematics units. It would be high risk to move them all at once to an AUTOSAR methodology and standard, but there are wide ranging benefits for both OEMs and Tier 1 suppliers in making the transition. It is estimated that by 2020, all vehicles will have some AUTOSAR-based ECUs, so this standard cannot be ignored.

Some of the reasons and benefits for switching to AUTOSAR include:

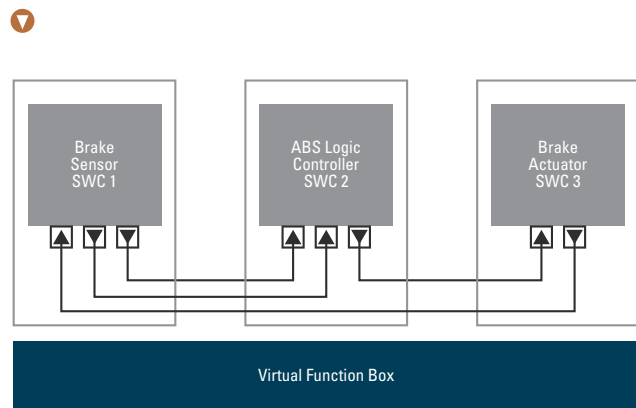
- Improved reuse of ECUs in new car platforms and architectures
- Improved use of pre-validated and tested software components (representing vehicle functions)
- Reduced testing and safety certification costs
- Reduction in downstream design errors—an AUTOSAR methodology allows functions to be defined and verified at an architectural level
- Reduction in overall hardware cost by improved network efficiency and capacity utilization
- Reduced costs in overall network architecture analysis and design reviews
- Improved communication between OEMs and Tier 1 suppliers, by using a standardized digital exchange format (AUTOSAR XML or arxml)

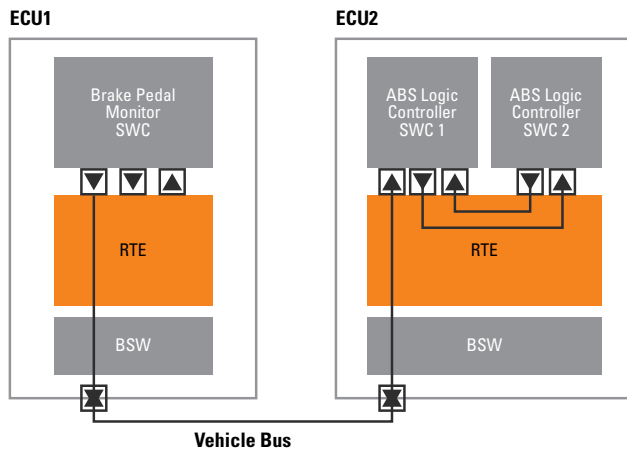
The move to AUTOSAR can be used as a catalyst for change—whether it’s an ECU that needs to be redesigned, or improvements required in the overall in-house design cycle. A move to an AUTOSAR methodology can be introduced alongside other process changes, such as a new tooling workflow, or adoption of improved safety standards to help with ISO26262 conformance. However the change is implemented, the first AUTOSAR-based ECU design project will take longer than the existing/legacy design process, as engineers become familiar with the new methods. The cost savings and efficiency benefits will follow later. It is also possible to migrate legacy ECU assets to AUTOSAR—using the concept of an “AUTOSAR wrapper,” in this way valuable existing and proven ECU application code can be reused. The AUTOSAR-enabling wrapper is capable of interfacing to other pure AUTOSAR ECUs.

At its heart, AUTOSAR provides a standard ECU interface definition, and allows an engineer to specify standardized reusable software layers and components that need to exist in every automotive ECU. The standard is hardware-independent, which means a clear line is drawn between the application software and hardware platform. The application developer can specify the details of individual vehicle functions in the application software without worrying about underlying software services and hardware interfaces. In the past, software and hardware had been very tightly integrated, hindering portability and reuse (Figure 1).

Abstracting the design away from hardware decisions opens up a freedom for top-down design by the vehicle manufacturer/OEM based on the required functions of the vehicle. For this stage of the design process, the concept of a Virtual Function Bus (VFB) exists, which allows all the software ECUs to be interconnected and tested. In addition, by using a VFB, the application software components (SWCs) do not need to know about other application software components. The software components present their output to the supporting VFB, which passes the information to the input ports of the destination components. AUTOSAR defines the input and output ports as well as the format of information exchanged. This abstracted approach makes it possible to validate the interaction of all vehicle software functions and interfaces

**Figure 2**  
Testing software components (SWCs) on a Virtual Function Bus (VFB).





**Figure 3**  
Assigning software functions into actual ECUs.

before defining underlying hardware. It is also much easier to make design changes while all functions are defined as software elements on a VFB (Figure 2).

The VFB has no knowledge of how ECUs will later be distributed and interconnected in the real vehicle, but is a very useful testbench for the architectural design phase. Timing checks can

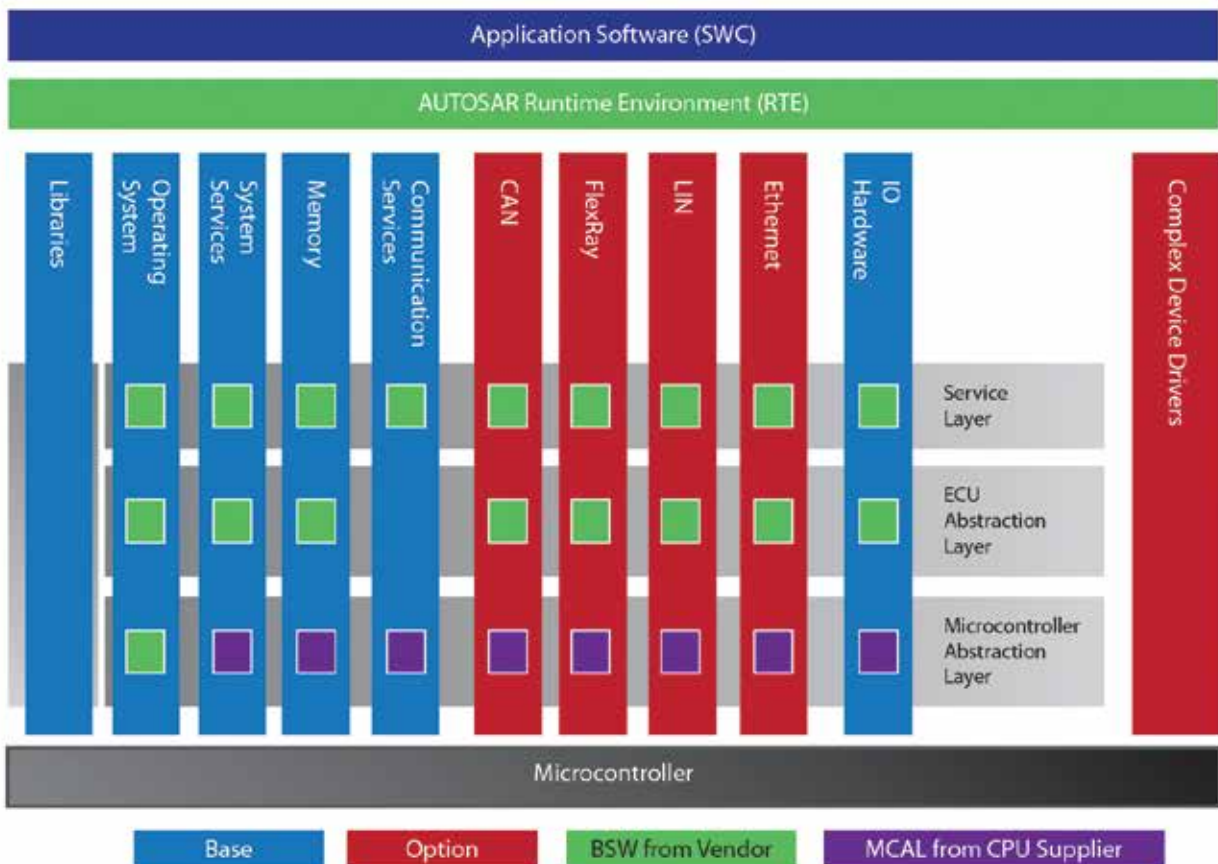
be carried out and the interfaces defined for all the signals in the vehicle.

Once the designer is satisfied with the individual functions, the functions are mapped or grouped together into specific hardware ECUs. AUTOSAR supports the process of mapping and grouping of software components (SWCs). A complex ECU may contain many SWCs, which can be organized hierarchically if necessary (Figure 3).

### The AUTOSAR Run-Time Environment (RTE)

Each individual ECU has its own customized run-time environment (RTE) implementation, which is normally automatically created by supporting design tools. The actual communication between real ECUs will be realized as part of a CAN or FlexRay bus, for example, and the RTE is configured by the generating tool to implement the communication paths required by its connected AUTOSAR components. The RTE becomes the “real-life” implementation of the communication and connectivity topologies from the VFB and architectural design process. Since the AUTOSAR standard supports many different types of software components, the

**Figure 4**  
How the components fit together in a real ECU.



RTE implementation must take into account the constraints and variations that a wide range of SWCs will introduce.

The Basic Software (BSW) is standardized software that does not contain vehicle application logic and ECU functions, but offers hardware-dependent and hardware-independent services to the RTE that sits on top of it. Examples of the services needed are memory services (NVRAM Manager), network communication management services, diagnostic services and state management. When an AUTOSAR SWC defined in the application layer requests services, it is the task of the RTE to complete the mapping on an actual ECU.

The RTE does not provide any mechanisms to access a service from a remote ECU, and this is not allowed by the AUTOSAR specification. All service requirements need to be fulfilled on the "local" ECU. The underlying operating system (OS or OSEK) that executes on the real ECU does not know about the concepts of AUTOSAR "runnables." The operating system maintains a list of schedulable events that are under management of a scheduling algorithm.

The AUTOSAR layered software architecture decouples the application logic from the hardware to facilitate reuse and portability. The RTE and operating system interface to the Microcontroller Abstraction Layer (MCAL), which in turn provides access to the physical ports and devices on the host microcontroller. The MCAL is specific to each microcontroller and allows the operating system and BSW to have access to devices such as digital I/O, analog-to-digital conversion, FLASH and EEPROM support, etc. Figure 4 shows the relationship between the different hardware and software layers in an AUTOSAR ECU.

### Enabling a New Methodology

With a top-down AUTOSAR design methodology, the automotive OEM can work with a complete model of the entire network. AUTOSAR design tools allow an individual ECU to be extracted, and the connectivity and interface information is defined in AUTOSAR XML (arxml). This interface definition can then be passed to a Tier 1 supplier for further detailing and implementation. Because the format is standardized, the same definition can be passed to several Tier 1s at the time of competitive tender. The standard description has the benefit of avoiding any design ambiguities in the ECU description, and as the AUTOSAR standard evolves, there is ever less room for misinterpretation. The standard is already hardware-independent so it is well placed to take advantage of new industry trends, such as Ethernet in the vehicle, mixed technology vehicle networks (CAN/Flexray), heterogeneous multicore platforms and in-vehicle gateway arrangements.

Several commercial organizations, including Mentor Graphics, offer evaluation kits for AUTOSAR design. These kits cover both architectural designs down to the configuration of individual ECUs. Mentor Graphics also has its VSX tool suite, as well as ECU hardware development boards with CAN, FlexRay, LIN and Ethernet support. The tools are Eclipse-based and make use of open source tool chains to take designs from source code through to run-time implementation. A low risk investigation and trial of AUTOSAR is preferable to a "big-bang" approach where ECUs in a vehicle are migrated all at once to an AUTOSAR methodology.

The AUTOSAR standard brings with it the opportunity for process improvement and component reuse, but also the challenge of learning a new ECU design process and tooling. The early adopters of AUTOSAR have been passing this knowledge into mainstream engineering, and resources and production-ready tools are now widely available. The adoption of AUTOSAR is also helping organizations meet their requirements of functional safety standard ISO26262, as it provides for a repeatable, well-defined, top-down design process.

**Mentor Graphics**  
**Wilsonville, OR.**  
**(503) 685-7000**  
**www.mentor.com**



## Breaking the Chains!

### Open and Flexible System Architecture for Safe Train Control

Rugged Computer Boards and Systems for Harsh, Mobile and Mission-Critical Environments

- Modular, SIL 4-certifiable systems for safety-critical railway applications
- Configurable to the final application from single function to main control system
- Communication via real-time Ethernet
- Connection to any railway fieldbus type like CANopen, MVB, PROFINET, etc.
- Comes with complete certification package including hardware, safe operation system and software
- Compliant with EN 50155



[www.menmicro.com/markets/railways.html](http://www.menmicro.com/markets/railways.html)





## Multi-Level Cell SSDs perform in HPEC Rugged Environment

Cost, performance and reliability are often conflicting goals.

by Steve Gudknecht and Ken Grob, Elma Electronic

Because all solid state flash drive (SSD) products are not created equal—and because flash storage is practically a given for use in high performance embedded computing (HPEC) applications—system designers should understand the critical tradeoffs between competing flash technologies when evaluating SSDs.

The high data rates required in constantly evolving electronic systems require extremely reliable performance to ensure data integrity. With more systems being deployed in mobile and harsh environments, operation across an extended temperature range is now more than ever a crucial factor.

Major enhancements in NAND flash controllers—the front end traffic cop of every SSD sold—can cost-effectively ensure the necessary performance in these harsher environments. Most commonly, the endurance and reliability required in end-user applications help dictate which storage technology is the most appropriate to use. Two well-known NAND flash storage technologies, single level cell (SLC) and multi-level cell (MLC), offer distinct advantages and tradeoffs depending upon a user's needs.

### Technology Differences: MLC vs. SLC

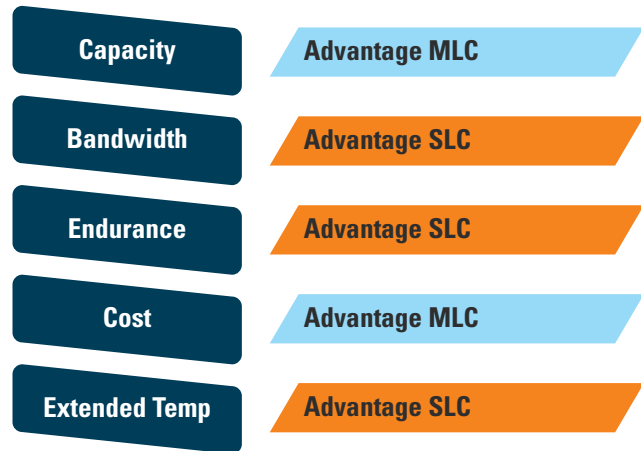
The difference between MLC and SLC is in the voltage level treatment and control within the cells. SLC stores a single bit in two binary states in each cell, while MLC stores two bits in four binary states in each cell.

In either case, the binary states are determined by the differentiation in charge levels on the floating gate. Advances in MLC technology are turning out 3- and 4-bit chips with eight and 16 binary states per cell, respectively, that allows higher densities per unit area in MLC-based drives and hence higher capacities in a given drive form factor.

The downside is that increasing the number of potential binary states within a cell reduces the delta between voltage thresholds, blurring the distinction between cell values, especially in the face of cell degradation over the life of the drive.

At higher temperatures, cell degradation is accelerated, and this thinner cell value separation is why MLC flash is more susceptible to higher bit error rates at extended temperatures. Higher bit error rates push the need for controllers with tighter and more time-consuming error correction algorithms in MLC, leading to slower read/write rates compared to SLC.

In a nutshell, many challenging embedded applications often demand conflicting requirements (Figure 1):



**Figure 1**  
Both MLC & SLC have distinct advantages, depending on the application needs.

- High capacity – advantage: MLC at +40% in leading edge capacity
- High bandwidth – advantage: SLC almost twofold
- High endurance – advantage: SLC by a factor of 20
- Low cost – advantage: MLC at 70% below SLC
- Extended temperature operation – advantage: SLC at equal bandwidth

### Market Need

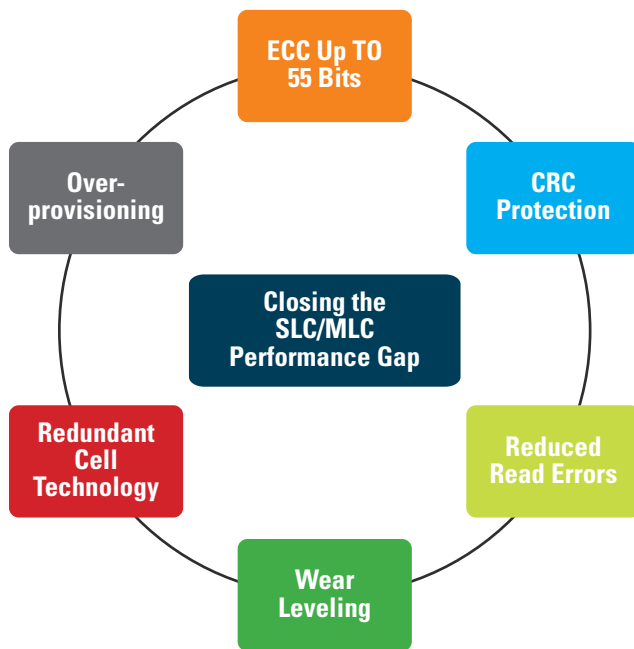
Low price and higher drive capacities in MLC are attractive to cost-sensitive industrial and mil program designers who

# TRACE32® Trace-based Code Coverage

**Real-time  
No instrumentation**

**LAUTERBACH**  
DEVELOPMENT TOOLS

[www.lauterbach.com/1659](http://www.lauterbach.com/1659)



**Figure 2**

Advancements in NAND flash controllers have lessened the performance gap between MLC & SLC.

consider deploying them in harsh environments, where hot and cold extremes are common and endurance is paramount.

In an age of ever-tightening budgets, economic viability must be achievable across platform designs and cannot be a barrier to implementation, but when SLC drives cost on the order of 2 to 4 times that of equivalent capacity MLC drives, you can guess where the money wants to go.

Lower MLC bandwidth issues can be addressed via improved flash controllers inside the drives, however it would seem that trouble will arise when low cost/high capacity requirements clash with high endurance/high temperature requirements. Looks can be deceiving and unlike the saying, “Get it good, fast and cheap...pick any two,” the need to compromise is fast diminishing as advanced SSD controller technology helps boost MLC performance.

## Extending MLC into the SLC Space

Until recently, MLC flash was only recommended for use in commercial temperature environments where the top end approached only 70°C. Many military applications, as well as some industrial applications, call for rugged equipment to operate at temperatures up to 85°C. While SLC solutions are certainly able to address those applications from an environmental and operational standpoint, the sticking point has always been the entry price.

Increased operational temperatures have the nasty habit of causing higher bit error rates and so SSD suppliers enlist a barrage of controller-based management features designed to mitigate if not eliminate the effects of bit errors. Major SSD manufacturers like Soligen, Memkor and others provide industrial temperature validated MLC drives. So advancements in NAND flash controllers have begun to close the MLC/SLC gap and SSD manufacturers, according to Mark Ayers of Soligen, are taking full advantage of that fact. Ayers notes that modern NAND controllers offer features such as (Figure 2):

- Error correction code (ECC) recovery of up to 55 bits correctable per 512-byte sector
- Unrecoverable read errors of less than 1 sector per 1017 bits read
- ECC on all internal memory and end-to-end cyclic redundancy check (CRC) protection
- Redundant NAND cell technology providing additional levels of data protection
- Over-provisioning for enhanced drive life span
- Wear leveling for enhanced drive endurance and life span

As you can imagine, all of that ECC overhead takes its toll on bandwidth and is the reason why inherent MLC bandwidth at the chip level is about 30% lower compared to SLC. Help has arrived, however, in the form of more effective internal flash organization, including I/O lane multiplication and improved management processes resident in the controller firmware.

These internal enhancements increase the number of read/write operations per second and reduce the lower bandwidth issue for MLC. In real terms this simply means that, thanks to flash controller technology and internal architecture, a 2.5” SATA 3 drive will operate at the specified theoretical maximum bandwidth of 6 Gbit/s, regardless of whether it’s MLC or SLC. The same goes for all form factor drives. Taken in that absolute sense, MLC and SLC drives of identical form factor can be equal

## Example Storage Recorder Topology

OpenVPX, or VITA 65, supports HPEC recording and is built on VITA 46 using the multi-gig connector that supports differential signals such as SATA or SAS used in building storage arrays.

An 8-channel OpenVPX RAID Controller (Figure 3) packages two 2 1/2” SSDs per 3U VPX card. Up to eight drives can be mounted across four slots. This array uses one drive controller/dual drive carrier card and three OpenVPX dual drive carrier cards. The current capacity is 8 Tbyte, and will soon be extended to 16 Tbyte as drive capacities increase.

Typical data rates can reach as high as 380 Gbit/s, depending on other variables such as operating system and driver optimization, and the controller can operate successfully in -40° to +85°C environments.



in bandwidth capability, even though their underlying cell technology is not the same.

Many applications are well suited for MLC drive technology since the write/erase rate is inherently low. As the aggregate ingest bandwidth increases in today's data gathering applications, the system back-end must respond. In platforms that must buffer, store or perform a combination of preprocessing and buffering, the system record rate is driven to increase and capacities must rise.

According to Ayers, single mission data collection applications—where large blocks of sequential data are recorded at a single pass—are perfect environments for MLC drives. In many recording applications for video data, radar data and flight data logging, the data recording capacity exceeds the maximum single mission flight time. In these cases, the recorded data is erased after each mission download. MLC can safely provide many years of service in military and industrial environments depending on the traffic usage case.


Modern MLC or SLC NAND-based SSDs have already achieved very good read and large block sequential write performance in these types of applications, according to Wieslaw Wojtczak, CTO of Memkor. The key focus for SSD manufacturers, comments Wojtczak, is to enhance small block random write performance by minimizing various latencies and in high speed multi-SSD RAID applications—by minimizing RAID response delays.

According to Wojtczak, in addition to long used DRAM caching, the most efficient approach so far in enhancing SSD bandwidth (MLC or SLC) is the use of overprovisioning. Overprovisioning, in this case, sets aside as much as 20-30% of the drive capacity as a work space for smart data processing and allows flash management background processes, like garbage collection, to run more effectively. Flash storage uniquely mandates that all old data must be erased before new data can be written. Since the entire flash block is erased at one time, valid data must be separated from invalid and moved to another location, so the block with only invalid data can be erased. This is known as garbage collection.

This overhead write penalty is a key factor in “write amplification” as it slows down write speeds and increases cell degradation. Smart data processing and optimized garbage collection routines allow increased performance, ensure performance stability over the usage of the SSD and decrease the write amplification factor. Overprovisioning is most important for enterprise class applications, where users access data at an extremely high rate with a high constant rate of random writes.

## Endurance and drive life

Wider design and operating margins in SLC flash lead to higher reliability, better endurance and longer life when compared to MLC flash. Drive manufacturers specify endurance as the number of block level write/erase cycles allowed before errors rise to unacceptable levels causing drive failures. SLC flash endurance stands at 10 to 20 times that of MLC.



**Figure 3**  
Open VPX RAID controller.

To address whether to use SLC or MLC, systems that rely on SSDs need a reliable method to alert users of the useful remaining drive life and then generate warning messages intended to trigger preventative maintenance. This process, known as self-monitoring, analysis and reporting technology—or SMART—is an open standard algorithm that enables SSDs to continuously monitor cell wear and send an alert to the user of impending cell failures.

To maximize drive life, wear leveling is employed. Wear leveling uniformly spreads writes over the entire physical array rather than continuously writing information to the same blocks. The controller manages this by maintaining a virtual map of the flash surface that points to the ever changing physical locations.

As blocks reach their endurance maximum, they are relegated to bad block status and no longer accessed. Wear leveling and SMART are key features when considering any type of flash, but are particularly important where MLC flash is concerned.

## Conclusion

Widespread adoption of MLC flash-based SSDs in military and industrial applications is being enabled by advancements in flash controller technology. Operation in extreme temperatures imposes performance demands that test the inherent limitations of MLC flash, as designers seek to take advantage of its low cost in these applications. Controller-based flash management technologies, like wear leveling, error correction, garbage collection, overprovisioning and SMART, combine to make MLC flash a worthy contender for usage in appropriate harsh environment applications.

**Elma Electronic**  
**Fremont, CA.**  
**(510) 656-3400**  
**www.elma.com**

# RTC PRODUCT SHOWCASE



**Networking Platform with Atom C2000 and Robust Security**

A 1U rackmount hardware platform designed for network service applications supports the next generation Intel Atom Processor C2000 product family (codenamed Rangeley). The choice of three high performance, low-power SoCs yielding 1.70, 2.0, and 2.40 GHz in dual- and quad-core packages is expandable to include other Atom C2000 processors with special orders. Robust security features include support for Intel AES New Instructions (Intel AES-NI), Intel QuickAssist Technology and Intel Streaming SIMD Extension (Intel SSE) for hardware accelerated data encryption and decryption.

Intel QuickAssist technology provides hardware level cryptographic acceleration. Off-loading compute-intensive security tasks provides additional processing power for higher layer packet processing by the CPU. The PL-80610 from WIN Enterprises is a suitable platform for system integrators (SI), electronic OEMs and software developers that provide networking and network security solutions the Enterprise, SMBs, and remote/satellite offices. Typical applications include network intrusion prevention and detection, unified threat management (UTM), spyware control, and content filtering.

The PL-80610 supports Intel 22nm Atom C2000 processors (codename Rangeley and is OEM-customizable to support other Intel Atom C2000 processors, including 8-core packages. The platform supports four DDR3 1333/1600MHz unbuffered ECC or non-ECC DIMM sockets up to 16GB of memory. The device supports 2.5"/3.5" SATA 3.0 6Gbps hard drives and Compact-Flash. A flexible 8 GbE to 15 GbE Ethernet ports are provided via PCIe on the front-panel. To prevent network problems during any unexpected shut downs, PL-80610 supports three segments of LAN bypass function through WDT and GPIO pin definitions. For local system management, maintenance and diagnostics; the front panel is equipped with dual USB 2.0 ports, one RJ-45 console port and LED indicators to monitor power and storage device activities. The PL-80610 has a PCIe X8 slot to support a range of Ethernet expansion modules.

**WIN Enterprises North Andover, MA  
(978) 688-2000. [www.win-ent.com](http://www.win-ent.com)**



**App Enables Remote Access of Modbus TCP I/O Modules**

A new app brings the power of mobile connectivity to industrial data acquisition and control. The free app Sealevel Modbus Connect for iOS 7 and later from Sealevel Systems allows easy communication with Sealevel Modbus TCP compatible products including SeaI/O and eI/O modules. Use Sealevel Modbus Connect to access the registers, coils and discretes of your Modbus device from your connected iPhone or iPad. The app includes low-level Modbus support that separates the hardware data acquisition layer from the software application layer, simplifying set up and configuration.

Key product features include the ability to monitor and control Modbus TCP I/O from iPhone or iPad, the ability to graphically display I/O status and to easily Read and Write coils, discrete inputs, input registers and holding registers. In addition, the user can view Modbus TCP raw frames and send custom Modbus commands.

Plant engineers can use Sealevel Modbus Connect to monitor and control I/O remotely, saving time and decreasing downtime. I/O status is displayed using easy to understand graphical screens that support both portrait and landscape orientations.

The app is a powerful tool for design engineers and programmers developing Modbus applications. The in-depth access to Modbus operation quickly identifies errors that might otherwise require extensive troubleshooting. Users can easily send custom Modbus requests and access raw request and response frames. The app can even be used as an educational tool for anyone that wants to learn Modbus.

**Sealevel Systems, Liberty SC. (864) 843-4343  
[www.sealevel.com](http://www.sealevel.com)**



**High Performance PXI Express Controller with 4th Gen Core i7**

A high performance 3U PXI Express (PXIe) embedded controller is equipped with the quad-core fourth generation Intel Core i7-4700EQ processor and operates at up to 3.4 GHz (in single-core Turbo Boost Mode). With four links x4 or two links x16 and x8 PCI Express Gen 2 link capability, up to 8 GByte/s of total system throughput and up to 16 GB of DDR3L 1600 MHz RAM, the PXIe-3985 is suitable for applications requiring intensive data analysis or processing and high-speed data streaming, such as in wireless, radar, or RF testing environments.

Based on PCI Express technology, the Adlink PXIe-3985 can offer four links x4 or two links x16 and x8 PXI Express link capability for interfacing with a PXI Express chassis backplane. When configured in this combination, with a high performance PXI Express chassis such as the Adlink 3U 18-slot PXES-2780 chassis, maximum system throughput of up to 8 GByte/s is enabled, providing an effective solution for high bandwidth applications requiring intensive data analysis or processing and data streaming. In addition, the PXIe-3985 features up to 16 GB of 1600 MHz DDR3L memory capacity, ideal for seamless execution in memory-intensive applications.

The Adlink PXIe-3985 provides versatile I/O capability, including dual DisplayPort connectors with 4K 2K support, dual GbE ports, GPIB, four USB 2.0 ports, dual hi-speed USB 3.0 ports, and trigger I/O for advanced PXI trigger functions. Easy connection with external, standalone instruments or devices, and innovative designs—such as dual BIOS backup—reduce maintenance efforts, with fast and easy swapping of battery, storage device, and SODIMM modules to deliver high availability in testing systems.

The PXIe-3985, supporting Windows 7 32/64-bit operating systems, provides optimal performance when installed in the Adlink 3U high capacity 18-slot PXES-2780 chassis or compact 9-slot PXES-2590 chassis. These combinations offer an ideal operating environment for a wide variety of testing and measurement applications.

**IBASE Technology, Taipei, Taiwan,  
886-2-26557588 [www.ibase.com.tw](http://www.ibase.com.tw)**



**3.6 GS/sec Rugged Portable RF/IF Signal Recorder**

A rugged portable recorder, suitable for military and aerospace applications is equipped with a 3.6 GHz 12-bit A/D converter. The RTR 2729A from Pentek is capable of capturing an extremely wide band of signals in real time to disk. The user-programmable digital downconverter (DDC) allows the system to capture tunable IF signals with bandwidths up to 360 MHz continuously for over four hours. The RTR 2729A is based on a new packaging scheme that boasts a smaller package, lighter weight and faster data rates.

The RTR 2729A uses a Pentek Virtex-7-based Onyx software radio board with a PCIe Gen. 3 engine to provide data streaming for the high-speed A/D converter. Coupled with a high-performance PCIe Gen. 3 SATA III RAID controller, the RTR 2729A is capable of streaming contiguous data to disk in real time at rates up to 4.8 GB/sec, which is 2.4 times faster than the previous generation.

The RTR 2729A features a portable, lightweight housing measuring only 16.0" W x 6.9" D x 13.0" H, weighing just less than 30 pounds. This extremely rugged workstation is reinforced with shock absorbing rubber corners and an impact-resistant protective glass for its high resolution 17" LCD monitor.

The hot-swappable Solid State Drive (SSD) array is available in 7.6 TB and 15.3 TB configurations and supports RAID levels 0, 1, 5, or 6. The SSDs are meticulously qualified by Pentek for optimum use in rugged and portable applications. The hot swappable solid-state drives exhibit high immunity to shock and vibration for full operation in ground vehicles, ships and aircraft. Available I/O includes audio and VGA video, six USB 2.0 ports, two USB 3.0 ports and dual Gigabit Ethernet connections.

All Talon RTR Portable Recorders are built on a Windows 7 Professional workstation with an Intel Core i7 processor and provide both a GUI (graphical user interface) and API (Application Programmer's Interface) to control the system. Systems are fully supported with Pentek's SystemFlow® software for system control and turn-key operation. The software provides a GUI with point-and-click configuration management and can store custom configurations for single-click setup. The software also includes a virtual oscilloscope and signal analyzer to monitor signals before, during and after data collection. The Talon RTR Portable Recorders start at \$74,995 USD. Delivery is 6-8 weeks ARO for all models.

**Pentek, Upper Saddle River, NJ  
(201) 818-5900. [www.pentek.com](http://www.pentek.com)**



# ADVERTISER INDEX



## GET CONNECTED WITH INTELLIGENT SYSTEMS SOURCE AND PURCHASABLE SOLUTIONS NOW

Intelligent Systems Source is a new resource that gives you the power to compare, review and even purchase embedded computing products intelligently. To help you research SBCs, SOMs, COMs, Systems, or I/O boards, the Intelligent Systems Source website provides products, articles, and whitepapers from industry leading manufacturers—and it's even connected to the top 5 distributors. Go to Intelligent Systems Source now so you can start to locate, compare, and purchase the correct product for your needs.

 [www.intelligentsystemssource.com](http://www.intelligentsystemssource.com)

Company .....	Page .....	Website .....
congatec, Inc .....	4 .....	<a href="http://www.congatec.us">www.congatec.us</a>
Dolphin .....	5 .....	<a href="http://www.dolphinics.com">www.dolphinics.com</a>
General Micro Systems, Inc.....	15 .....	<a href="http://www.gms4sbc.com">www.gms4sbc.com</a>
Lauterbach Development Tools.....	37 .....	<a href="http://www.lauterbach.com">www.lauterbach.com</a>
Men Micro, Inc.....	35 .....	<a href="http://www.menmicro.com">www.menmicro.com</a>
MSC Embedded Inc .....	4 .....	<a href="http://www.msembedded.com">www.msembedded.com</a>
One Stop Systems .....	7,19 .....	<a href="http://www.onestopsystems.com">www.onestopsystems.com</a>
Pentek, Inc. ....	2 .....	<a href="http://www.pentek.com">www.pentek.com</a>
Portwell.....	44 .....	<a href="http://www.portwell.com">www.portwell.com</a>
Real-Time & Embedded Computing Conference.....	42 .....	<a href="http://www.rtecc.com">www.rtecc.com</a>
Trenton Systems.....	43 .....	<a href="http://www.trentonsystems.com">www.trentonsystems.com</a>
TQ Systems GmbH.....	25 .....	<a href="http://www.convergencepromotions.com/TQ-USA">www.convergencepromotions.com/TQ-USA</a>
WDL Systems .....	11.....	<a href="http://www.wdlsystems.com">www.wdlsystems.com</a>

RTC (Issn#1092-1524) magazine is published monthly at 905 Calle Amanecer, Ste. 250, San Clemente, CA 92673. Periodical postage paid at San Clemente and at additional mailing offices.  
 POSTMASTER: Send address changes to The RTC Group, 905 Calle Amanecer, Ste. 250, San Clemente, CA 92673.

## The Event for Embedded & High-Tech Technology

---

2014 Real-Time & Embedded Computing Conferences


Tysons Corner Area, VA – November 18

Santa Clara, Ca – January 22

Register today at [www.rtecc.com](http://www.rtecc.com)



# Trenton's THD8141 System Host Board Has It All



- > Four Standard DDR3-1600 Plug-in DIMMs
- > Multiple Video Ports
- > SATA/600 RAID
- > USB3 Ports
- > Latest Multi-Core Processor Options
- > Standard PCI Express 3.0 and PCI Plug-in Card Support



Trenton's THD8141 is our latest system host board to feature support for systems utilizing both standard off-the-shelf PCI Express® 3.0 and PCI plug-in cards. This new PICMG® 1.3 SHB offers a choice of multi-core Intel® Xeon® E3-1200 v3 or an Intel® Core™ i3/i5/i7 processor. The THD8141 system host board also includes:

- The Intel® C226 Platform Controller Hub (PCH)
- Three Gigabit Ethernet interfaces
- Revision Control BIOS & BIOS Customization Services

The THD8141 has all you need for a wide variety of embedded computing systems. Longevity and performance requirements are covered in military computing applications. THD8141 enables system design flexibility in telecom, video command & control and energy exploration. Other board advantages include:

- PCI Express link, device I/O and media expansion options
- 22nm Intel® Micro-Architecture (i.e. Haswell) performance
- Five-year warranty & seven years of product availability

*Our board design experts are available to discuss your unique application requirements.  
Contact us to learn more at 770.287.3100 / 800.875.6031 or [www.TrentonSystems.com](http://www.TrentonSystems.com)*

The Global Leader In Customer Driven Computing Solutions™

770.287.3100 | [www.TrentonSystems.com](http://www.TrentonSystems.com) | 800.875.6031






# Portwell Empowers Intelligent Solutions



Mini-ITX



Small Form Factor System



COM Express Module



Network Security Appliance



PICMG SBC



[www.portwell.com](http://www.portwell.com)  
[info@portwell.com](mailto:info@portwell.com)  
1-877-278-8899

