INCLUDING DEFENSE TECH WIRE

Military EMBEDDED SYSTEMS

April/May 2012
Volume 8 | Number 3
MIL-EMBEDDED.COM

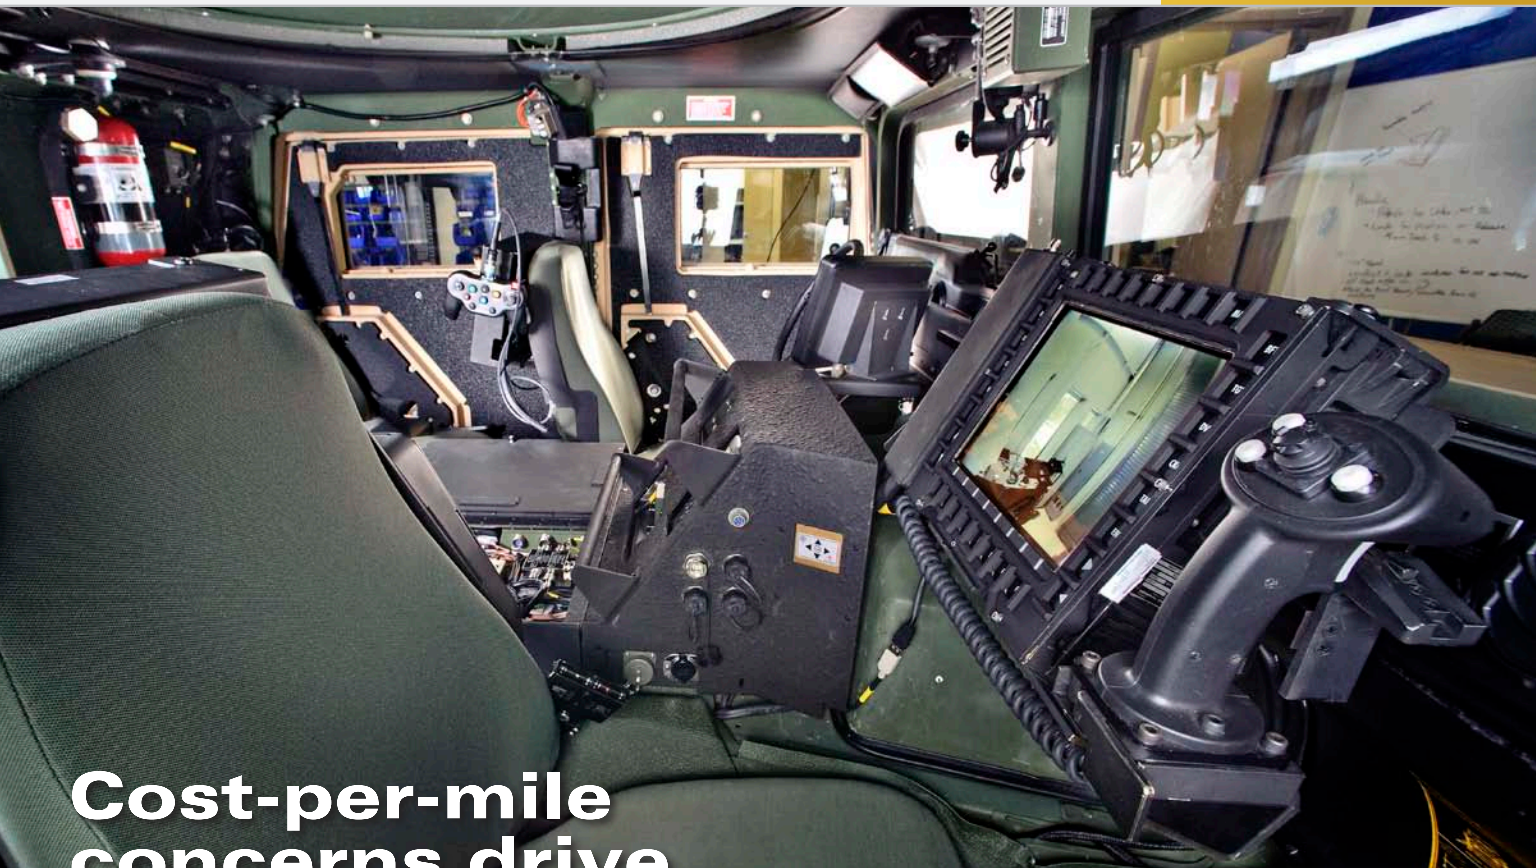

Cost-per-mile concerns drive military vetronics designs

Secure virtualization for tactical environments by David Egts, Red Hat

SBCs: Increased performance, lower power, business as usual

ELECTRONIC SERVICE REQUESTED
16626 E AVENUE OF THE FOUNTAINS, STE. 201, FOUNTAIN HILLS, AZ 85268
PRST STD
U.S. POSTAGE
PAID
OpenSystems Media

Military EMBEDDED SYSTEMS

www.mil-embedded.com

ON THE COVER:

Top photo: Pictured is an interior view of a four-seat Up-Armored M1165 HMMWV. The vetronics system in the HMMWV was designed by General Dynamics C4 Systems, Vehicle C4ISR Innovation Center in Taunton, MA. Photo courtesy of General Dynamics C4 Systems.

Bottom photo: An EA-18G Growler assigned to the Shadowhawks of Electronic Attack Squadron 141 prepares to take off from the aircraft carrier USS George H.W. Bush. U.S. Navy photo by Seaman Billy Ho

April/May 2012 | Volume 8 | Number 3

COLUMNS

DEPARTMENTS

EVENTS

www.opensystemsmedia.com/events

Aerocon

May 2-3, 2012 • Charlotte, NC
www.aeroconshows.com

Military, Defence & Aerospace Electronics Technical Conference & Exhibition 2012

May 22, 2012 • Reading, UK
www.mae-show.com

WEB RESOURCES

Subscribe to the magazine or E-letter
Live industry news | Submit new products
http://submit.opensystemsmedia.com

White papers:
Read: http://whitepapers.opensystemsmedia.com
Submit: http://submit.opensystemsmedia.com

Published by: OpenSystems media

ISSN: Print 1557-3222

16

26

34

@military_cots

www.linkedin.com/groups/Military-Embedded-Systems-1864255

ADVERTISER INFORMATION

Military EMBEDDED SYSTEMS

OpenSystems media.

Military Embedded Systems Editorial/Production Staff

John McHale, Editorial Director
jmchale@opensystemsmedia.com

Sharon Hess, Managing Editor
sharon_hess@opensystemsmedia.com

Terri Thorson, Senior Editor (columns)
tthorson@opensystemsmedia.com

Steph Sweet, Creative Director
ssweet@opensystemsmedia.com

Sales Group

Dennis Doyle
Senior Account Manager
ddoyle@opensystemsmedia.com

Tom Varcie
Senior Account Manager
tvarcie@opensystemsmedia.com

Rebecca Barker
Strategic Account Manager
rbarker@opensystemsmedia.com

Eric Henry
Strategic Account Manager
ehenry@opensystemsmedia.com

Ann Jesse
Strategic Account Manager
ajesse@opensystemsmedia.com

Christine Long
Director of Online Development
clong@opensystemsmedia.com

International Sales
Elvi Lee, Account Manager – Asia
elvi@aceforum.com.tw

Regional Sales Managers
Barbara Quinlan, Midwest/Southwest
bquinlan@opensystemsmedia.com

Denis Seger, Southern California
dseger@opensystemsmedia.com

Sydele Starr, Northern California
sstarr@opensystemsmedia.com

Ron Taylor, East Coast/Mid Atlantic
rtaylor@opensystemsmedia.com

Reprints and PDFs

republish@opensystemsmedia.com

OpenSystems Media Editorial/Production Staff

CompactPCI AdvancedTCA & MicroTCA | DSP-FPGA.com | Embedded Computing Design | Industrial Embedded Systems | Military Embedded Systems | PC/104 and Small Form Factors | VITA Technologies

Mike Demler, Editorial Director
DSP-FPGA.com
mdemler@opensystemsmedia.com

Joe Pavlat, Editorial Director
CompactPCI, AdvancedTCA, & MicroTCA Systems
jpavlat@opensystemsmedia.com

Jerry Gipper, Editorial Director
VITA Technologies
jgipper@opensystemsmedia.com

Warren Webb, Editorial Director
Embedded Computing Design
Industrial Embedded Systems
wwebb@opensystemsmedia.com

Jennifer Hesse, Managing Editor
Embedded Computing Design
Industrial Embedded Systems
jhesse@opensystemsmedia.com

Sharon Hess, Managing Editor
VITA Technologies
sharon_hess@opensystemsmedia.com

Monique DeVoe, Assistant Managing Editor
PC/104 and Small Form Factors
DSP-FPGA.com
mdevoe@opensystemsmedia.com

Brandon Lewis, Associate Editor
CompactPCI, AdvancedTCA, & MicroTCA Systems
blewis@opensystemsmedia.com

Curt Schwaderer, Technology Editor

Steph Sweet, Creative Director

David Diomede, Art Director

Joann Toth, Senior Designer

Konrad Witte, Senior Web Developer

Matt Jones, Web Developer

Editorial/Business Office

Patrick Hopper, Publisher
Tel: 586-415-6500
phopper@opensystemsmedia.com

Subscriptions
Karen Layman, Business Manager
www.opensystemsmedia.com/subscriptions
Tel: 586-415-6500 ■ Fax: 586-415-4882
30233 Jefferson, St. Clair Shores, MI 48082

Rosemary Kristoff, President
rkristoff@opensystemsmedia.com

Wayne Kristoff, CTO

16626 E. Avenue of the Fountains, Ste. 201
Fountain Hills, AZ 85268
Tel: 480-967-5581 ■ Fax: 480-837-6466

March embedded shows, VPX, code analysis

By John McHale, Editorial Director

"In like a lion, out like a lamb" is an oft-used cliché when describing the third month of the year, but it works quite well as a description of my March 2012 trade-show schedule, which began in Nuremberg, Germany and ended in San Jose, CA.

The Germans roared into March with Embedded World, an annual gathering of everyone and every gadget involved in embedded technology. Long acknowledged as the inventors of the trade show more than 1,000 years ago, German organizers have not lost their touch. This year's show drew more than 20,000 attendees and nearly 900 exhibitors. The huge crowds were steady through the end of the show.

Attendees in San Jose at the Design West show – a.k.a. the Embedded Systems Conference (ESC) – did not make as much noise, but seemed just as energized. Renamed *Design West* for 2012, the event had steady traffic, even on the last day. The majority of attendees and exhibitors I spoke with felt they had a good event filled with strong meetings, even though overall attendance was modest compared to Embedded World earlier in the month.

" 'In like a lion, out like a lamb' is an oft-used cliché when describing the third month of the year, but it works quite well as a description of my March 2012 ..."

Neither event demands pilgrimage from military program managers, but there were a few military-related product and standards developments discussed at both shows.

For example, VPX was notable for the interest it sparked at each event, but especially in Germany. Though Embedded World caters to the industrial market, more than a few Single Board Computer (SBC) vendors there said many military integrators inquired about VPX-based products. Also, two of the top five viewed stories posted on our website during Embedded World were about VPX. One was an item about a new OpenVPX board from Eurotech, and the other was on VPX interest in Europe increasing. For more on our coverage of both events, please visit www.mil-embedded.com. For more on VPX and VME trends, see our Mil Tech Trends articles beginning on page 22.

However, even with the increased military interest in embedded technology, Embedded World will always be mainly an industrial event. For that matter, it is unlikely that military trade shows in Germany will ever be a common occurrence, as the German companies that support military programs are still quite gun-shy about broadcasting those efforts. Obviously, this cultural shyness stems from World War II. As one German resident told me, the German people fear any perception that they are beating war drums and want their nation to be seen as a welcoming, hospitable place for foreign visitors. They succeeded in the latter.

The Europeans also want Europe to become a more hospitable place for standards development. Frustrated by what they say are long standards development processes within certain U.S. organizations such as PICMG, leaders of different European embedded computing vendors formed a new standards consortium called the *Standardization Group for Embedded Technologies* or *SGET,* which was announced during Embedded World. Founding member companies of SGET include Advantech, congatec, Data Modul, Kontron, MSC, and SECO.

My favorite product from both shows was a new tool from GrammaTech that actually makes code analysis fun. The new tool – which works with GrammaTech's CodeSonar static analysis tool – functions much like Google Earth, but instead of zooming in on your neighbor's house, you're zooming down to view individual lines of code in systems that contain as many as 20 million lines of code. The tool won the 2012 Embeddy Award for most significant new software product at Design West. Our post on the GrammaTech tool also was the most viewed article on our website during Design West.

Paul Anderson, Vice President of Engineering for GrammaTech, told me that the design team originally based the tool on Google Earth, but found that model inefficient and so they created their own visualization tool that still has the same "zooming" effect that you see with Google Earth. Users can actually scroll via the mouse all the way out so that the program is just a dot, like the Big Bang theory, then shoot all the way down to individual lines of code. Anderson said the visualization tool was designed for a customer program that contained 20 million lines of code, but there's no reason to think the tool cannot analyze even larger amounts of code.

GrammaTech grabbed quite a bit of attention with their award, but I think the most hardware collected during Design West was by an attendee who won the last giveaway at our booth on the show's final day. I was excited for him since he never missed a drawing, but that empathy faded when I spotted him outside resting his arms on a brand new 40" flat panel screen. Some people collect business cards, winners take home big TVs.

John McHale
jmchale@opensystemsmedia.com

Record/Playback @ Record Breaking Speed...

D-TA'S 10 GIGABIT SENSOR PROCESSING (10 GSP) implementation has revolutionized sensor processing by offering **limitless scalability** that allows the throughput rate to increase linearly as more networks (fibers) are added, and **totally synchronized operation** that ensures the data is precisely time synchronized across all fibers, irrespective of the number of fibers.

Record/Playback @ over 800 MBytes/s with 1 fiber

D-TA (or 3rd party product)

1 X 10 GbE Network

Compact 1U Solution

DTA-1000-R
(1U & 4 TB SSD Drives)
[8 TB with SATA]

Record/Playback @ over 4 GBytes/s with 4 fibers

D-TA's Sensor Interface Product (or 3rd party product)

4 X 10 GbE Networks (Fibers)

2 X DTA-5000
(6U & 19.2 TB SAS Drives)
[64 TB with SATA]

Record/Playback @ over 2 GBytes/s with 2 fibers

D-TA (or 3rd party product)

2 X 10 GbE Networks

Stackable 3U Solution

DTA-5000
(3U & 9.6 TB SAS Drives)
[32 TB with SATA]

We are delivering recording solutions that sustain over 8 GBytes/s using 8 fibers and 4 DTA-5000 units.

For more information, click the **Recording at Record Breaking Speed** button on our home page at **www.d-ta.com**

Leveraging our 10 GSP concept, we have developed a generic record / playback configuration that measures up to your requirements. Radio, radar, sonar, acoustic — whatever your application, we likely have an end-to-end solution from signal acquisition to record & playback.

D-TA SYSTEMS INC.
A Sensor Interface and Processing Company

10 GIGABIT SENSOR PROCESSING
FAST • SCALABLE • SYNCHRONIZED

SENSOR PROCESSORS THAT DRASTICALLY REDUCE DEPLOYMENT TIME AND COST

Crunch time for signal processing

By Charlotte Adams

In line with the new *"lean and mean"* U.S. defense strategy, the military is developing advanced, nearly autonomous sensor platforms. Collectors on Unmanned Air Vehicles (UAVs), for example, are expected to siphon oceans of raw data from the environment, process much of it onboard, and provide useful information in real time.

This combination of data wealth and processing constraints will require not only faster hardware but also more complex application code. Platforms are to read, process, and exploit data from wide-area, multispectral, or hyperspectral sensors, sifting out targets and points of interest within tactical timelines. Drones will linger over an area, tag day-to-day changes such as new encampments or improvised explosive devices, track local personnel movements, and perhaps identify specific individuals.

Tough order

Apart from the obvious technical challenges, developers face shortened program cycles because of military budget and operational imperatives. Meanwhile, the underlying platforms such as ships and aircraft are expected to endure for decades longer than before, making smooth software evolution and technology insertion as important as the original applications.

Programs such as Gorgon Stare and Autonomous Real-Time Ground Ubiquitous Surveillance-Imaging System (ARGUS-IS) are examples of efforts to provide wide area, persistent surveillance information in real time. ARGUS-IS reportedly uses a 2 gigapixel sensor with hundreds of focal plane arrays, producing images that must be stitched together to provide a continuous view.

Upgrade programs can be iterative, releasing incremental improvements at intervals as short as 12 months. Companies can also fund technology demonstrations to generate customer interest. Either route puts a premium on development speed. Software development has to keep pace.

Chip makers are rising to meet the challenge. One example is Intel's latest Core i7 product, code named *Ivy Bridge,* which is expected to hit the street about the same time as this column. Reportedly, Ivy Bridge will shrink the second-generation IC manufacturing process to 22 nm, while squeezing 1.4 billion transistors onto a smaller die. The mobile version of the product, appropriate for embedded applications, is expected to include four cores, each with its own 256-bit-wide vector processor, a feature introduced by the predecessor chip. Such a vector pipeline, by allowing multiple operations to execute simultaneously, would be well suited to signal processing, where key functions require repetitive operations across large data sets.

Figure 1 | DSP280 rugged, dual-socket, Intel second-generation quad Core i7 multiprocessor from GE Intelligent Platforms

Toolkits to the rescue

How will developers exploit such advances? Toolsets from hardware and software companies include large math libraries that relieve programmers from coding basic signal processing building blocks such as FFTs, matrix operations, and filter functions. Some toolkits also accelerate productivity by orchestrating the work of potentially hundreds of processors. Where symmetric multiprocessing is used, the software engineer can write a single program that will be automatically partitioned across all available cores.

Graphical tools help the developer to build an application by structuring it via manageable tasks and automatically generating code to distribute the software and initialize the system. Companion tools help depict central processing unit usage and communications channel bandwidths, as well as identify bottlenecks, diagnose bugs, and find timing discrepancies.

Toolkits are provided by multiple sources such as chip and board makers. GE Intelligent Platforms' Advanced Multiprocessor Integrated Software (AXIS) environment, for example, includes math libraries, visualization tools, and interprocessor communications optimizers that run on multiprocessors such as the company's DSP280, a rugged, dual-socket, Intel second-generation quad Core i7 multiprocessor (Figure 1).

Some tools also allow the programmer to simulate an application on a desktop computer before moving it to multiprocessor embedded hardware.

The most useful toolkits embrace multiple chip architectures and generations, as well as multiple operating systems, including the real-time operating systems powering embedded applications. These toolkits also allow the programmer to exploit all hardware features with no detailed knowledge of their inner workings.

Programming force multiplier

As chips become more capable and complex, software development tools are becoming more comprehensive yet easier to use, putting signal processing application development on a glide path to greater efficiency. Weapons developers as well as weapons platforms can eventually become force multipliers.

Contact Charlotte at cburtonadams@yahoo.com.

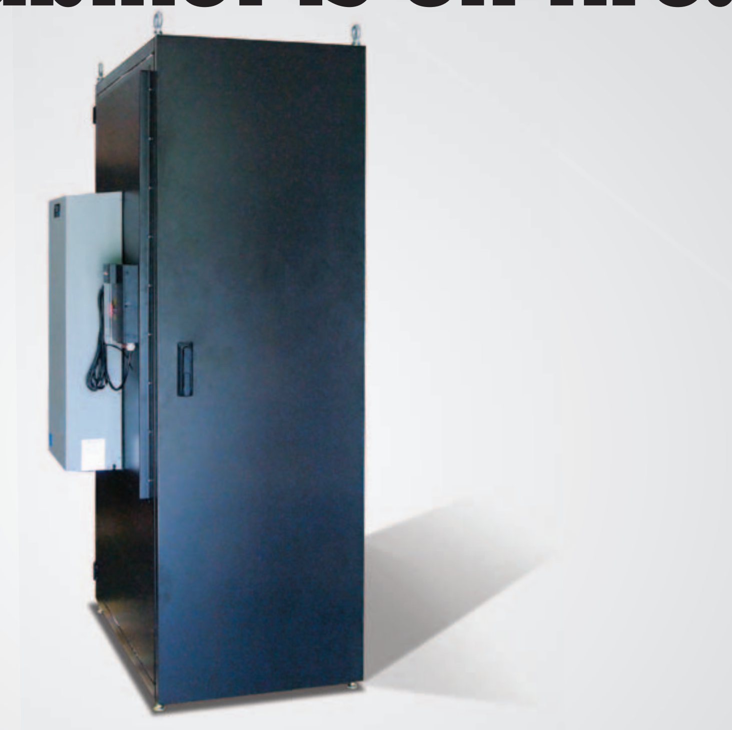

The VICTORY architecture drives COTS-based interoperability on combat vehicles

By Steve Edwards

Combat vehicles today typically have multiple independent systems that have no ability to share their functionalities or data. The U.S. Army's Vehicular Integration for C4ISR/EW Interoperability (VICTORY) initiative encourages the use of COTS open-system standards, reduces redundancy, and makes additional space available by reducing Size, Weight, and Power (SWaP). The VICTORY initiative, which is not tied to any specific program or platform, coordinates input between government, academia, and industry participants. The DoD's commitment to a Modular Open System Architecture (MOSA) is driving the use of open-system approaches both from the technical and the procurement perspectives.

The VICTORY initiative is an instantiation of the MOSA policy. The initiative's charter is to define a standard approach for intravehicle networking to drive interoperability and drastically reduce component redundancy and the resulting inefficient use of limited real estate in combat vehicles caused by the use of "stovepipe" or "bolt-on" subsystems. At its core, VICTORY is developing standards for interoperability between Line Replaceable Units (LRUs) on combat vehicles. VICTORY defines the use of nonproprietary interfaces between heterogeneous LRU subsystems. The resultant open architecture standard won't define how LRUs, typically supplied by different vendors, are built, but how they can intercommunicate and share data and resources.

Open, interconnected, interoperable environment

A useful metaphor for today's stovepipe combat vehicle LRUs is to consider what desktop computing would be like if one needed a separate stand-alone system, likely from different vendors, for word processing, spreadsheets, and presentation software. Essentially, this is the operating environment in today's combat vehicles, where there is limited interoperability and almost no sharing of data or resources. This is the exact problem that VICTORY aims to eliminate.

Because VICTORY LRUs communicate over distributed IP- based networks (exploiting cutting-edge commercial networking technology such as Web Services, SOAP, and XML), VICTORY helps eliminate redundant system components such as displays, keyboards, and GPS receivers. The use of open-system architectures ensures that customers have true choice with access to best-in-class, cost-effective technologies by eliminating the deployment of proprietary technologies. Further, by establishing a standard approach for LRU network interfaces, VICTORY eases and speeds the use of LRUs across various platforms, bringing the proven COTS model already embraced on the Line Replaceable Module (LRM) level to the LRU. Open standards and multiple vendor choices will lower costs by reducing the subsystem "lock-in" typical in today's larger programs.

Figure 1 | The SwitchBox SMS-685 Ethernet router intended for VICTORY compliance from Curtiss-Wright Controls Defense Solutions

For an embedded COTS vendor such as Curtiss-Wright Controls Defense Solutions (CWCDS), VICTORY means a level playing field to compete and provide the market with a full range of interoperable COTS-based C4ISR components including computers, network switches, application processors, and multiple LRUs for use on many vehicles. VICTORY ensures that products from competing vendors can interoperate, easing the customer's integration task while significantly reducing design risk and time to market, all to benefit the warfighter's effectiveness.

On the road to VICTORY 2.0

The VICTORY standardization process began in May 2010. An initial timeline was established, targeting release of a deployable, usable revision 1.0 VICTORY Specification by the June/July 2011 timeframe. The Revision A VICTORY Architecture standard, describing the approach, components, and services, was released in April 2011, and the 1.0 VICTORY Specification, detailing requirements and message sets, debuted June 2011, as scheduled. In early 2012, the VICTORY Specification 1.1 was released, adding more subsystem message sets and recommended standard military rugged (MIL-STD-38999) connectors for Ethernet connections, adopting a standard in use by the United Kingdom's Ministry of Defense Generic Vehicle Architecture (GVA). The VICTORY working groups will continue to incorporate additional vehicle and C4ISR/EW subsystems in the road map to VICTORY 2.0.

VICTORY compliance for new programs

Vendors are already beginning to see VICTORY compliance as a requirement for new light tactical vehicle subsystems as well as in the emerging modernization requirements for combat vehicle platforms. Open standard-based rugged subsystem suppliers like CWCDS can address critical programs such as Abrams, Bradley, Stryker Modernization, and the Joint Light Tactical Vehicle with technology and packaging leadership, delivering standard-compliant products and their resulting flexibility, modularity, compatibility, and reduced cost. Figure 1 is an example of an Ethernet Router intended for VICTORY compliance.

To learn more, e-mail Steve at Steve.Edwards@curtisswright.com.

DEFENSE TECH WIRE

mil-embedded.com/magazine/wire

By Sharon Hess, Managing Editor

NEWS | TRENDS | DoD SPENDS

Northrop Grumman to boost DoD cybersecurity

The number of cyberthreats has increased for virtually anyone who uses or owns a computer these days. Accordingly, the Defense Information Systems Agency (DISA) has extended a $189 million task order to Northrop Grumman to boost DoD and intelligence community cybersecurity (Figure 1). Additionally equipped with two one-year options, the task order specifies that Northrop Grumman will render architectural, maintenance, engineering, deployment, and integration support for the DoD-mandated Host Based Security System (HBSS). HBSS is COTS software utilized by the DoD to accelerate laptop- and desktop-based host security, rather than focusing on boundary security of switches and routers. The goal is to thwart not only external, but internal, cyberthreats. The contract covers field activities and agencies, Services, and combatant commanders, as well as the intelligence community's host platforms and networks, among others.

Figure 1 | DISA has awarded a $189 million task order to Northrop Grumman to boost DoD and intelligence community cybersecurity. U.S. Army photo by Spc. Loren Cook

ISR ordered for Special Ops

It's not just military embedded industry talk. The U.S. Special Operations Command is indeed using Intelligence, Surveillance, and Reconnaissance (ISR) technology. Case in point: AAI Corp. of Hunt Valley, MD was granted a three-year, up to $600 million contract for ISR services for U.S. Special Operations Command pertaining to the Mid-Endurance Unmanned Aircraft System II. Under the contract, the company will render contractor-operated and -owned equipment. The contracting activity is the U.S. Special Operations Command Headquarters Procurement Division at MacDill Air Force Base in Florida.

L-3 night vision goggles aid Special Ops

U.S. Special Operations Command missions are bound to be in the dark sometimes, where all the intelligence and weapons in the world won't make a difference – if operatives can't see what's ahead. Thus, the Naval Surface Warfare Center has awarded a $50 million IDIQ contract to L-3 Communications for Binocular Night Vision Devices (BNVDs) that will be utilized by the U.S. Special Operations Command. The BNVD goggles are geared to render increased detection, recognition, identification, and surveillance capability, even in low-light situations. The goggles comprise a high-performance image intensifier and advanced optics in a mini, two-tube configuration. The BNVD is produced by L-3's Warrior Systems-Insight Technology unit.

Lockheed Martin instrument passes sensitivity test

The Geostationary Lightning Mapper (GLM) instrument – part of NASA and NOAA's Geostationary Operational Environmental Satellite (GOES)-R Series (Figure 2) – recently completed its sensitivity testing. Specifically, the new-to-GOES instrument successfully endured its optical-electronic lightning sensitivity exam, and will undergo thermal-vacuum and thermal testing in a few months. The near-infrared GLM comprises a 1.8 megapixel, 500 frames per second local plane, melded with specialized optics, with the goal of detecting faint lightning signals even in sunlight. GLM will map lightning over the Americas, as well as over nearby oceans.

Figure 2 | Lockheed Martin's Geostationary Lightning Mapper (GLM) instrument – part of NASA and NOAA's GOES-R – successfully completed its sensitivity testing. Logo courtesy of NASA.

For consideration in Defense Tech Wire, submit your press releases at http://submit.opensystemsmedia.com. Submission does not guarantee inclusion.

General Dynamics/U.S. Army's triad of contracts

When the U.S. Army and General Dynamics put pen to paper, they did so three times in one day, spanning two General Dynamics divisions and various technologies: 1) The first "cost" contract is $11 million, modifies an existing contract, and stipulates that General Dynamics Land Systems renders long-lead materials for M1A2S tank production. [The M1A2S is the Saudi Arabian version of the M1A2 tank (Figure 3.)] Contract completion is anticipated this November. 2) Next, a General Dynamics Land Systems/U.S. Army contract for nearly $8 million, again as a contract modification, specifies that General Dynamics drives efforts for M1A2S tank conversions. Contract work is expected to be completed by this May. 3) And finally, a third contract, yet again serving as a contract modification, stipulates that General Dynamics Armament and Technical Products will convert the machine guns dubbed "M2" into an M2A1 configuration. The contract is slated for completion by May 2013.

Figure 3 | A triad of General Dynamics/U.S. Army contracts covers M1A2S tank production/conversions plus M2-to-M2A1 transformation. M1A2 photo by Sgt. Quentin Johnson, 2nd BCT PAO, 1st Cav. Div., U.S. Army

Boeing/USAF contract: If looks could kill

In the case of the Boeing-incarnated Joint Helmet Mounted Cueing System (JHMCS), the old adage "if looks could kill" is not so far-fetched. Specifically, a recent USAF/The Boeing Co. contract provides for $31 million worth of full-rate production for the JHMCS, which, when paired with a mini display system, enables JHMCS's magnetic helmet-mounted tracker to become a targeting device capable of aiming weapons and sensors at the location where the pilot is looking, according to the Boeing website. Work under the contract covers 111 systems to be deployed to the U.S. Navy, along with Australia, Belgium, Pakistan, Finland, Switzerland, and Canada. Contract completion is slated for December 2013.

Figure 4 | A $19.5 million U.S. Navy/Raytheon Co. AIM-9X missile contract modification covers the needs of 10 countries. U.S. Navy photo by Mass Communication Specialist 3rd Class Christopher K. Hwang

Raytheon contract mod serves 10 countries

A recent U.S. Navy/Raytheon Co. AIM-9X Sidewinder missile contract modification carrying a $19.5 million price tag spans the needs of 10 countries (Figure 4). Specifically, the AIM-9X integrated logistics support contract "buyers" are the U.S. Navy (32.98%) and USAF (47.4%), in addition to the governments of Denmark, Australia, Singapore, South Korea, Turkey, Finland, Poland, Saudi Arabia, and Switzerland, each of which contributed 2.18% of the total remuneration. The majority of the work will be completed at Raytheon's Tucson, AZ locale; the remainder of the contract work is slated for fulfillment in other U.S. and international locations. The contracting activity is the Naval Air Systems Command in Patuxent River, MD.

E-2C trainers to be rejuvenated

The E-2C (Figure 5) simulated maintenance trainer and weapon system trainer will soon be spruced up with anti-obsolescence upgrades and software modifications, courtesy of a U.S. Navy/Rockwell Collins Simulation & Training Solutions, LLC, contract. Work under the $37.5 million contract will transpire primarily in Norfolk, VA (84 percent) and additionally in Point Mugu, CA (the other 16 percent). Contract fulfillment is slated for March 2017, and the contracting activity is the Naval Air Warfare Center Training Systems Division located in Orlando, FL.

Figure 5 | Rockwell Collins will soon be providing Navy E-2C (pictured) trainer upgrades and software mods. U.S. Navy photo by Mass Communication Specialist 2nd Class Walter M. Wayman

Lowering cost per mile in military vehicles drives vetronics designs

By John McHale, Editorial Director

Designers of military vetronics computers, displays, and software say integration and low power are driving requirements for upgrades and new systems.

Pictured is an interior view of a four-seat Up-Armored M1165 HMMWV that uses a vetronics system from General Dynamics C4 Systems that leverages an Open Systems Software Framework and embedded small form factor hardware.

Nobody likes new car payments. Just as the average consumer wants to make his car or SUV last forever, DoD program managers want their ground vehicle platforms to perform for at least another decade. Maintaining and keeping these vehicles up to date with current technology will be expensive. Success will be measured by how well the cost per mile can be kept low enough without sacrificing capability or performance.

Cost pressures are really driving the innovation path in vetronics, says David Jedynak, Manager – Advanced Solutions at Curtiss-Wright Controls Defense Solutions in Charlotte, NC. Cost per mile is a way to measure total cost of ownership of a ground vehicle, including the vehicle electronics, he adds. The cost-per-mile metric takes into account maintenance costs, fuel costs, training costs, and operating costs. The vetronics are just one little part of that, but a very expensive part. So the DoD sees it as an area to reduce expenses.

"Today military ground vehicle integrators are looking at cost, commonality, size, and high Technical Readiness Levels (TRLs), Jedynak says. "For C4ISR [Command, Control, Communications, Computers, Intelligence, Surveillance, and Reconnaissance] solutions, performance is still paramount because those applications demand cutting-edge solutions. But when talking about vehicle electronics, the key drivers are cost and the ability to have common parts across multiple vehicles, and high TRL levels, which mean the solution is ready to go and a lower risk and there's little or no Non-Recurring Engineering (NRE). Those are the more important factors right now."

SWaP-C optimization

"Performance is important for high-end systems such as signals intelligence and battle command, but today, for vetronics it's all about Size, Weight, [Power, and Cost] (SWaP-C) reduction: cost, commonality, and open standards that don't require a lot of space are the driving requirements," Jedynak says.

Smart displays meet military vetronics demands for reduced SWaP and better security

Smart displays meet military vetronics demands for reduced SWaP and better security

Integrators of vetronics systems for military vehicles are moving toward Commercial Off-the-Shelf (COTS) smart displays for new platforms and vehicle retrofits, as they offer reductions in Size, Weight, and Power (SWaP) and are more cost effective to develop.

"For displays we're seeing everything going digital in the vehicle," says Robert Kopersiewich, Vice President of Embedded Graphics & Mission Critical Solutions at Presagis in Montreal. The types of displays being developed are integrating more and more data for the warfighter, for example, some bringing in external feeds such as live video that is overlaid on the display.

Program managers want high-end graphics capability in a rugged, compact package that ensures they can run any application from any screen in the vehicle, says Greg Downs, Vice President of Vetronics Systems for General Dynamics Canada in Ottawa, Ontario. In other words, the driver should be able to see what the gunner sees and vice versa without having to get up and go look at the other's display, he explains. Smart display technology will solve these problems and give full access to any crew member, Downs adds.

When the vetronics displays are not integrated and mission-critical "information is not fully accessible, you don't have real crew situational awareness in the vehicles," Downs says. GD Canada produces smart displays that leverage Intel processor chipsets for enhanced video processing recording and high-end graphics processing, he says. The high-end graphics also enable embedded training and mission rehearsal on the platform while it is in the field. Being able to train warfighters in the field on the actual platform is also a major cost saver, Downs adds (Figure 1).

The typical dimension of displays in armored fighting vehicles is a ten-inch diagonal and a twelve-inch for battle command applications, Downs says. "We are seeing a desire to have larger displays on certain programs such as some fourteen-inch displays. However, the reality of space constraints will push it back from to ten to twelve inches. We expect that will remain the most popular size."

Security is another important trend in vetronics display designs, Downs says. The DoD wants classified and unclassified information to be accessible on the same display, he continues. They are looking for solutions "to automate the transition between security levels" so they can be consolidated onto a display or a single device. That consolidation also enables SWaP savings with the classified and unclassified data accessible on one display or computer as opposed to two of each, Downs says. Fewer displays mean more room in the vehicle and reduced power consumption, he adds.

Standard HMIs enable portability in displays

Developing display software that provides seamless integration between different vetronics displays is enabled by standards in software, specifically Human Machine Interface (HMI) technology and the ARINC 661 standard. For vetronics it is crucial to [design] HMI in a way that is deterministic and well defined for mission-critical applications, Kopersiewich says.

"ARINC 661 brings standardization across HMIs and does it cost effectively," Kopersiewich continues. "When an OEM is building land vehicles, they want to be able to outsource the HMI design knowing that it uses a common standard that can interface with other components of the platform's vetronics as well as being scalable and compatible with future variants," he explains. "ARINC 661 defines how displays are designed and feed into different user applications. It provides a kind of spoke and node architecture."

Sidebar Figure 1 | The 10.4-inch SD7310 smart display from General Dynamics Canada leverages Intel Core 2 Duo processors for high-performance graphics processing in vetronics applications.

The trend that is important here is that "COTS tools make more sense as the complexity builds up," Kopersiewich says. They can do hand-coding with OpenGL and create displays easily. It makes a lot of sense to have COTS tools used for vetronics applications, as the "vetronics displays are getting more and more complex," he adds.

The Presagis VAPS XT software is a COTS HMI tool that uses an object-oriented architecture and OpenGL interface for supporting HMI applications for use on different types of displays and touch screens, Kopersiewich says. With a tool such as the Presagis VAPS XT, users can reconfigure a single operator station quickly for many different missions. It is easier and more cost effective to change the design code in OpenGL instead of starting from scratch each time you need to change a display, he adds.

Sidebar 1 | The need for improved security and reduced SWaP is met by smart displays in military vetronics. Additionally, display portability is enabled by standard HMIs.

"The desire is to buy, train with, and stock/spare common components."

"Minimizing SWaP is the key trend – everyone wants a smaller and lighter subsystem with equal to or greater than the processing capability they already have with their existing subsystems," says Mac Rothstein, Product Manager, Systems, GE Intelligent Platforms in Charlottesville, VA. "High-performance computing is also trending up as more and more data needs to be captured and analyzed to give warfighters a clearer and more detailed picture of the battlefield for air and ground vehicles. GE's IPS5100 and IPS511 subsystems provide a 360-degree view of a vehicle's surroundings," he adds.

"Thinking capability first and power last has been a common theme in ground vehicle systems," says Doug Patterson, Vice President of Marketing at Aitech Defense Systems in Chatsworth, CA. "They are putting counter-IED [Improvised Explosive Device] jammer systems onto HMMWVs [High Mobility Multipurpose Wheeled Vehicles] and guess what? The battery and alternator are not big enough to handle all that power pouring from all that capability into the box. It requires kilowatts of extra power, which comes at a cost."

"At the electronics level, all the new processors such as the multicores, Intel Core i7s, etc., are creating a rise in temperature and the heat has got to go somewhere," Patterson continues. "So military integrators are saying to embedded suppliers: 'help us.' They are pushing for innovation from the board and subsystems designers on this issue," he adds. Each platform takes a different approach such as using liquid-cooling or spray-cooling techniques inside the enclosure/chassis or box.

Vetronics boxes use up a fair number of watts and when you have a lot of watts in an electronics box, you need a lot of metal to cool it – and a lot of volume is needed for the cooling mechanism, Jedynak says. Therefore, if you reduce the power you reduce the intrinsic size and weight of box and also material costs, he continues. "If we can push lower power with great performance, we get a strong cost reduction and a reduction in total cost of ownership."

Demands for Line Replaceable Modules (LRMs) are increasing and "we are seeing 3U viewed as a good way to address SWaP and cost constraints," Jedynak says. "Today, vehicles are overburdened in terms of SWaP: Whether they are tactical vehicles or combat vehicles like the Bradley Fighting Vehicle, they are dramatically overburdened because of the addition of armoring elements over the recent years. Therefore, reducing the size of computing elements, reducing weight and power, is a huge imperative today. So obviously, cost goes right along with that. The trend is moving toward commonality and 2-Level Maintenance, which drive us toward a blade-type computer system, i.e., 3U VPX LRMs."

Commonality and standardization are behind the U.S. Army's Vehicular

Integration for C4ISR/EW Interoperability (VICTORY) initiative. It provides interoperability at the subsystem level. "VICTORY is to subsystems as OpenVPX is to the backplane," Jedynak says. OpenVPX enables standard card modules from multiple vendors to work together within a chassis and "VICTORY shows us how to compose vehicle subsystems" that can work with each other via common standards, he adds. For more on VICTORY see the Mil Tech Insider column on page 12.

There are upgrade programs right now in initial stages for Bradley and Abrams, Jedynak says. Curtiss-Wright is looking to respond to these opportunities in terms of 3U VPX and commonality as well as SWaP-C optimized devices with VICTORY compliance. "Another important program is the JLTV (Joint Light Tactical Vehicle), which is a very cost-constrained vehicle" that will require low-cost, reduced SWaP solutions, he adds.

System integrators also are thinking literally outside of the box when it comes to managing power, Patterson says. Along those lines "we've got an outside-the-box, novel approach for distributing all that computing by putting remote interfaces that are concentrated around the vehicle. What we learned from the cancelled Future Combat Systems (FCS) program was not a waste. We leveraged real-world knowledge and the technology to distribute the computer architectures within the vehicle so that you don't have to concentrate all that heat/power into one spot. Distributing all of that power around the vehicle makes more sense than concentrating it in one place.

"The distributed computing concept echoes work the Army has done with wearable computing configurations on individual warfighters," he continues. "On a soldier you can put 60 or 100 watts and spread it across the body and the soldier will never know it, through cell phones, radios, wrist tablets, etc. Soldiers carry cell phones now and are not even thinking about the heat [they're] generating." **MES**

"On a soldier you can put 60 or 100 watts and spread it across the body and the soldier will never know it, through cell phones, radios, wrist tablets, etc. Soldiers carry cell phones now and are not even thinking about the heat [they're] generating."

Multifunction I/O Boards

You select the functions. We save you space.

Proudly made in the USA.

High density COTS I/O solutions reduce board count, weight and power... enhancing overall system reliability and affordability.

- Mix-and-Match from over 20 different I/O functions on a single board
- MIL-STD-1553, CANBus, ARINC 429/575, RS232/422/485, GbE Switch; Discrete, Differential, TTL; A/D, D/A, Synchro, LVDT, RTD; Local Processing
- Continuous background BIT monitors health and status of each channel
- User configurable I/O parameters, per channel
- Available Bus Architectures include VME, VPX, cPCI and PCI
- Access and control board via Bus Architecture or Gigabit Ethernet
- Modified COTS solutions also available

COTS Solutions tailored to meet your system requirements.
Visit **www.naii.com** or call us at 631-567-1100 today.

North Atlantic Industries
Excellence in ALL we do

Embedded Boards | *Power Supplies* | *Instruments*

631-567-1100 • Fax: 631-567-1823 • www.naii.com • email: sales@naii.com

Nothing new: Military users demand more performance and low power from SBC designs

By John McHale, Editorial Director

The major trend in the military Single Board Computer (SBC) market continues to be the demand from system integrators for greater processing performance and lower size, weight, and power in SBC designs. Meanwhile, the venerable VMEbus finds a small resurgence as funding increases for military retrofit programs.

U.S. Marine Corps photo by Sgt. Jesse J. Johnson

Wars may end or extend, defense budgets may shrink or grow, and elected officials will come and go, but designers of military electronic systems will always want Single Board Computers (SBCs) with more performance and lower power in smaller and smaller designs.

There is "nothing really surprising" going on right now in the SBC market, says Rob Scidmore, President and Chief Executive Officer of Extreme Engineering Solutions (X-ES) in Middleton, WI. "As it has been historically, for new programs, military customers want the latest and greatest processors and as much overall performance as they can get."

"We haven't seen a lot of change over the last year," says Mike Slonosky, Product Marketing Manager for PowerPC SBCs at Curtiss-Wright Controls Defense Solutions in Charlotte, NC. "What we are seeing is a tremendous need for lower Size, Weight, and Power (SWaP), especially in regard to the power per MIPS metric, which we believe the upcoming third generation of the Intel processor will improve, but don't know for sure."

Platforms such as combat vehicles and Unmanned Aerial Vehicles (UAVs) are becoming smaller and smaller, but still have a veracious need for processing power, says Alan Baldus, Product Marketing Manager for Intel SBCs at Curtiss-Wright Controls Defense Solutions.

"The footprint that's available for processing requirements within these platforms has gone down, but somewhat surprisingly, as the processing footprint has gone down, the processing requirements have gone up," Slonosky says. "We are seeing continual increases in demand for processing power with a simultaneous reduction in SWaP budget (including power and ability to remove heat). As a result, we are seeing a big increase in interest in 3U form factors." The Curtiss-Wright VPX3-1256 device uses the latest-generation Intel processor and is designed to track the Intel road map as it develops (Figure 1).

"We are seeing a trend toward smaller SBCs – fewer 6U boards and more 3U boards," Scidmore says. "We also are seeing increasing demand for smaller systems, which require the SBC to be even smaller than 3U. X-ES provides PrPMCs, processor XMCs, and

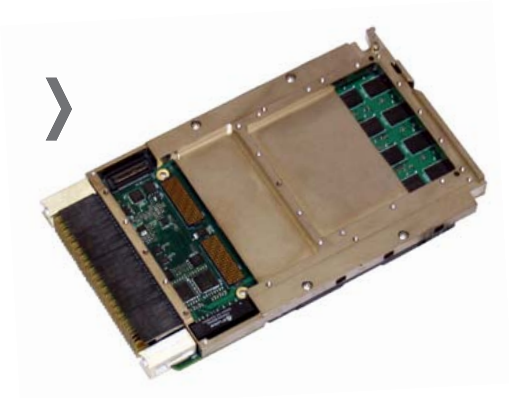

Figure 1 | The VPX3-1256 from Curtiss-Wright Controls Defense Solutions supports Quad SERDES and Dual DVI and can accommodate future Intel processors.

Figure 2 | The XPedite5550 is Extreme Engineering Solutions' first COM Express module.

just introduced its first COM Express module, the XPedite5550 (Figure 2). For more information, visit www.xes-inc.com.

Space consolidation within platforms is one of the major drivers of processing performance within military systems, Baldus says. "For example, in a ground combat vehicle with limited space, where today you might have five different independent systems and five different computers and five different screens, the desire going forward is to consolidate them as much as possible. If you can consolidate three of the five computers and display the resulting information on one common monitor, the benefit is a significant reduction in systems and weight, etc. in the vehicle."

"The types of systems that are in demand typically have four things in common: one, they need to be able to move large amounts of streaming data; two, they need to be able to process this data in real time; three, they need to be able to operate in harsh environments; and four, they need to keep the technology from being compromised by the enemy," Scidmore says.

Military users are looking for a "stable, reliable, proven building block on which to base the rest of the application, whether that application sees the integration of additional hardware or additional software," says Richard Kirk, Product Manager, Military/Aerospace SBCs, GE Intelligent Platforms in Charlottesville, VA. "Implementing the latest technology isn't just about performance, though – it's also about longevity. The earlier in the life cycle of new technology our customers can get it into their programs, the longer the solution can remain viable – and so the cost of ownership becomes lower. Fabric choice is also important to designers when integrating high-performance systems.

VME gets new life, while demand for VPX increases

VMEbus-based products continue to thrive, even as VME's "grandchild" VPX sees a large increase in demand and new products.

VPX will not eclipse VME everywhere, says Ray Alderman, Executive Director of the VITA Standards Organization in Fountain Hills, AZ (www.vita.com). "There are still VME niches in missiles and gun turrets, which need VMEbus for real-time responses. Systems with packet-based architectures will stay VME." Future enhancements to advanced systems will leverage VME-based designs in a big way, he adds.

VME is seeing a resurgence because of redeployment of existing military platforms, because the government doesn't have the money to do a whole new program, says Eric Gauthier, Vice President of Marketing, Strategic Alliances, and Partnerships for Embedded Computing at Emerson Network Power in Phoenix, AZ. Emerson's latest VME offering is the MVME8100, which uses a Freescale P5020 QorIQ processor. For more information, visit www.emersonnetworkpower.com.

"We continue to see demand for VME, and in some cases – programs that were looking to go VPX are staying VME," Slonosky says. With VME there is still continued interest, especially with the decrease in budgets, he continues. System upgrades to VPX will be slowed down and only used where it is most necessary, he says. Technology insertions for VME are still going strong too, as it does not always make sense to replace VME with VPX, Slonosky explains. For super high-performance applications, a move to VPX would be necessary, but VME is still an ideal form factor for many defense electronics applications.

DoD budget cuts are extending the life of existing programs for VME and CompactPCI for the next few years, says Norbert Hauser, Executive Vice President of Marketing at Kontron. This is good as both form factors are "cash cows for us." In the long term with new designs, VPX will be the standard of choice, he adds. Kontron plans to host the third-generation Intel Core processors on 3U and 6U CompactPCI processor boards, a 3U VPX processor board, and on other form factors. For more information, visit www.kontron.com.

The board-level market can be broken down into three levels – legacy VME in the middle with VPX at the high end and Small Form Factor (SFF) designs at the low end, Alderman says.

SFF designs will make use of ARM and Android technology rather than Intel chips, he continues. At the high end, VPX will be the choice, with processing technology from Intel and Freescale. Military applications using VPX will be radar, sonar, electronic warfare, signals intelligence, and mission computing, Alderman says. "There is a tremendous interest in radar for VPX." There may also be a bit of 3U CompactPCI being used at the high end, but the majority will be VPX, he says.

"In some recent discussions with some customers in Europe, we found that there is still interest in CompactPCI, especially in low-performance, low-cost systems, where for some applications VPX is considered too costly," Baldus says.

"For new design-ins, we're seeing quite a bit of interest in VPX both in 3U and 6U versions," says Ken Grob at Elma in PA. "We're also seeing requests for SBCs with heavy DSP capabilities that make use of FPGAs." Some designs also use GPGPUs for front-end DSP functionality, he adds. For more information on Elma embedded SBC products, visit www.elma.com.

Frank Godulla, Director of Sales for Hartmann Electronic, says VPX will really explode when interest peaks for VPX products outside of the military market. Right now the interest is mainly military, which means low-volume orders, he says. Once commercial-oriented markets integrate VPX systems, volumes will increase. The growth potential of VPX was also partially behind Hartmann Electronic buying a piece of MILPER, an Israel-based company that specializes in rugged computers and ATR systems, Godulla continues. Military customers who want VPX often want it in a rugged ATR system, and MILPER's technology enables Hartmann Electronic to provide those features, he adds. For more information, visit www.hartmann-electronic.com.

"VPX will really explode when interest peaks for VPX products outside of the military market."

Small Form Factors (SFFs)

The VITA 75 working group (RSFF or Rugged Small Form Factor) has produced a number of design concepts and draft specifications, Slonosky says. "We are diligently working on SFF subsystems through chairmanship of the VITA 75 working group, which addresses new SFF subsystems." For more information on CWCDS SFF products, visit www.cwcdefense.com.

The major applications for SFF designs are in vehicular electronics (vetronics); Intelligence, Surveillance, and Reconnaissance (ISR); and UAVs, says Dennis Smith, Vice President of Engineering at Themis Computer in Fremont, CA. Persistent surveillance is another area where SFF devices can thrive, he continues. They can be placed on robots or in "intelligent cubes" on top of telephone poles and provide wired 60-degree surveillance via "fish-eye lenses." Themis also is working with the infrared sensor community and other electro-optical sensing designers because "we see high-flying image processing as a way to reduce power and increase speed through small form factor designs."

The Themis VITA 74-based product – the NanoPak – has 5 slots and comes in 12.5 mm and 19 mm modules. The NanoPak is rugged and can be liquid and conduction cooled, Smith adds. For more information, visit www.themis.com.

Not all SFF designs are based on VITA 74 specifications as other SBC designers are developing products outside that specification.

"Some SFF efforts in the industry limit the processing power to processors such as the Intel Atom processor because of the inability to get heat out of the system," Scidmore says. "Some other companies base their SFF systems strictly around 3U modules, which significantly increases the size and reduces the performance capabilities of the system. With our SFF approach, the die of the processor is mated directly against the inside of the chassis' wall, which is either mounted onto a cold plate in a conduction-cooled application or has fins to remove the heat via convection. This allows us to cool and support these high-performance processors in very small systems."

"There are parts of our mil/aero business where small form factor is incredibly important, but in SBCs we're not seeing so much of a rapid trend to miniaturization," Kirk says. "Yes, for sure, managing SWaP continues to be a key concern for our customers. But, for the majority of these customers, 3U VPX gives them what they need, especially as the processors at the heart of those boards get more and more capable." GE Intelligent Platforms is "sponsoring the VITA 75 Rugged Small Form Factor standard, because that seems to us to be the best of the various competing standards that are around at present," Kirk says.

"Our attitude is colored by the premise that small form factor and standardization may not really go hand in hand – and the reality is that most customers wanting a small form factor solution will actually need a custom solution, because of the very specific constraints of each application." The latest product from the GE Intelligent Platforms VPXcel6 range of 6U VPX SBCs is the SBC624 (Figure 3), which uses the newest Intel processor and may be updated again in the near future, Kirk says. For more information, visit www.ge-ip.com. **MES**

Figure 3 | The SBC624 6U VPX board from GE Intelligent Platforms uses the newest Intel processor.

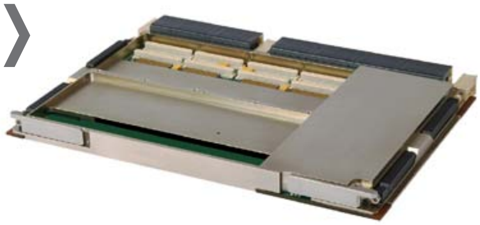

SBC TRENDS (PROCESSORS)

High-performance SBC form-factor trend forges military electronics

By Brad Trent

There are always engineering challenges related to designing and deploying single board computers for use in high-performance military systems. However, the increasing number of form factors on the market contributes greatly to improved compute density, and processor power and capabilities. An examination of two dual-processor, multicore SBC platforms as they are deployed on a military vehicle-mounted system and airborne surveillance platform in a PICMG 1.3 form factor exemplifies one such trend.

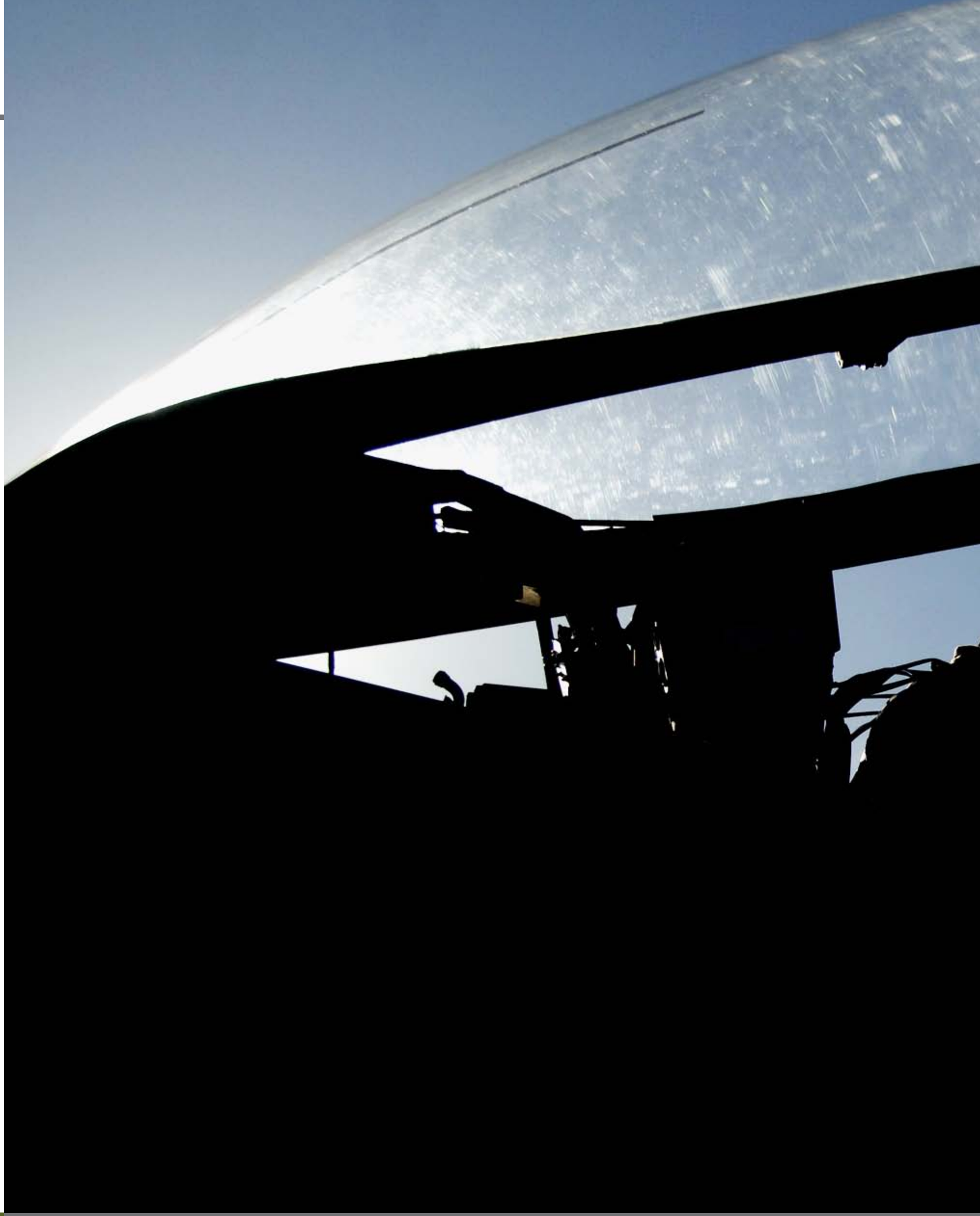

U.S. Air Force photo/Senior Airman Timothy Taylor

Single board computers are experiencing a renaissance in all sorts of military computing applications. Much of this can be attributed to the increasing number of different SBC form factors available to the military computer system designer: PICMG 1.3, PC/104, COM Express, MicroTCA, and 3U VPX boards, to name only a few. These SBCs offer inherent system architecture flexibility, while lending themselves well to advances in compute density and processor power and capabilities. Additionally, a couple examples of how to successfully deploy high-performance SBCs in military vehicle and airborne surveillance computing applications in a PICMG 1.3 form factor reinforce this trend.

SBC form factor and compute density

One of the first problems is to decide which form factor makes the most sense in the application. For example, if the SBC is targeted for a very specific series of tasks, then a smaller SBC form factor such as VPX, MicroTCA, or COM Express, for example, might be suitable with a scaled-down RTOS likely fitting the bill. One of the key system design advantages of these SBC form factors is that low-power processors can be used in order to design fanless computer boxes that can be embedded in the small spaces typically found in the cockpit of a fighter aircraft or under the seat of a ground assault vehicle.

But what about compute density? The smaller form factors are great for deploying in tight spaces and in applications with clearly defined operational tasks. However, where they fall short is when more compute capability is needed for tasks that may not be as clearly defined and include processing data from a wide variety of sources and communicating the resultant data analysis to various worldwide locations. An airborne surveillance aircraft, a shipboard navigation system, and government agency cryptographic threat analysis platforms are good examples of applications that need a large amount of computer processing capability housed in space-efficient platforms that maximize compute density. Figure 1 illustrates the differences between a small form factor SBC for a targeted application and a larger dual-processor SBC used in applications requiring processing capability and enhanced compute density.

A dual-processor, single board computer form factor with multiple PCI Express

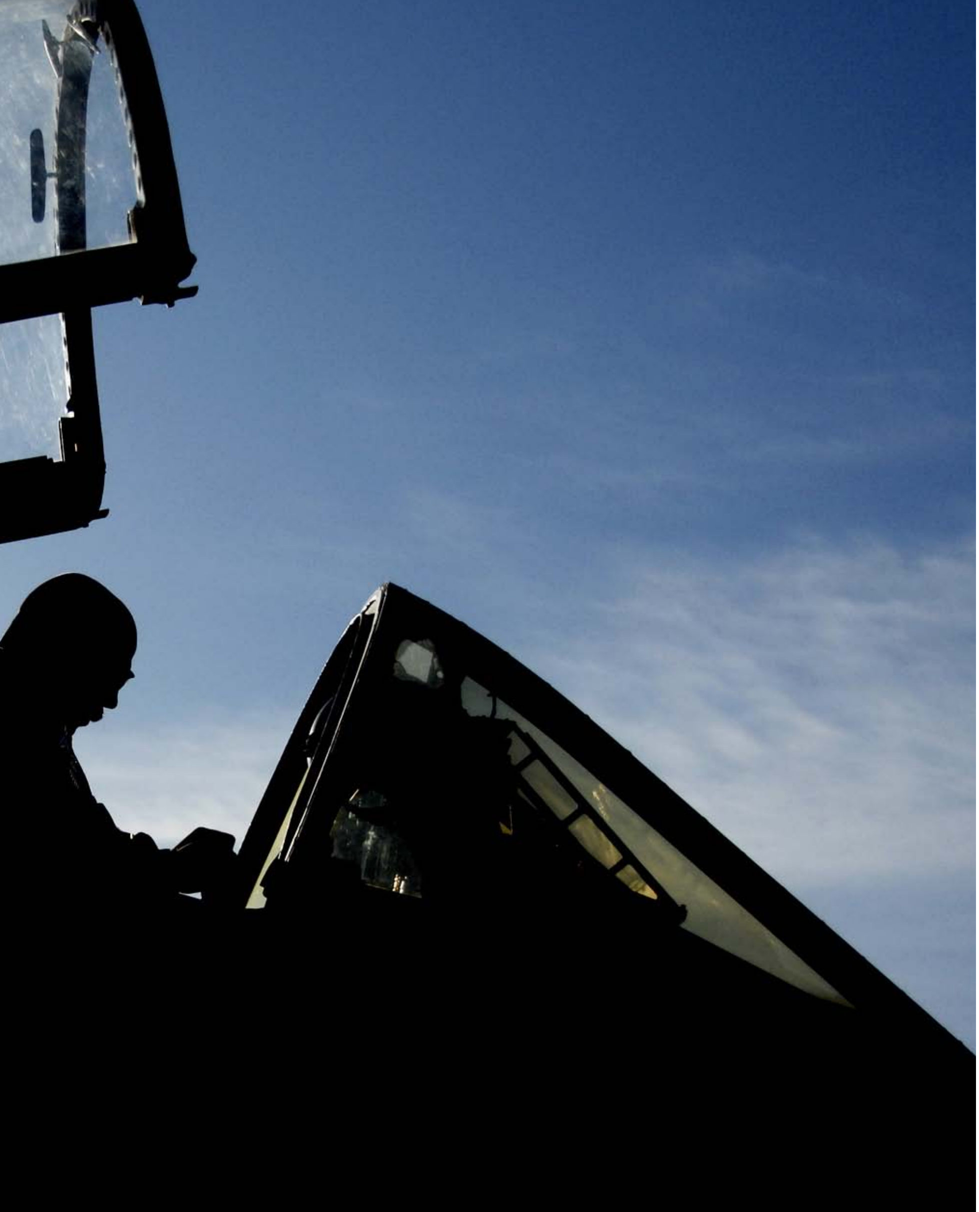

and Ethernet connections such as that typically found in PICMG 1.3 (SHB Express) System Host Boards or *SHBs* make PICMG 1.3 an ideal SBC form factor. It aptly addresses the need for military computer system performance and compute density. The following PICMG 1.3 specification highlights make this form factor especially useful:

- Sizes – Half-size = 6.6" x 4.98"/ 167.6 mm x 126.4 mm (max); full-size = 13.3" x 4.98"/ 338.66 mm x 126.4 mm (max).
- PCI Express – Support for up to 20 off-board lanes on PCI Express, which can be augmented with plug-in mezzanine cards. The PCIe lanes can be configured in a wide variety of configurations.
- PCI – The spec allows for an additional PCI parallel interface connection using an optional edge connector D on a full-length PICMG 1.3 SBC.
- I/O – Standard off-board interfaces such as USB, SATA, and Ethernet are defined on the edge connectors.
- UP or DP design support – A half-size PICMG 1.3 usually supports a single processor (UP) design while a full-length board can support either a single or Dual Processor (DP) design.
- System flexibility – Defined edge connector connection plus backplane form factors defined in the specification enable military system designers to mix and match SBCs and backplanes from multiple vendors in order to develop the best system solution possible.
- Compute density – The PICMG 1.3 form factor dimensions and edge connector/backplane approach enable rack-mount computer system designs utilizing standard caseworks that can offer multiple SBCs in a single enclosure.

See the PICMG website at www.picmg.org/v2internal/SHB_Express.htm for more information on the PICMG 1.3 specification.

There are certain engineering design techniques needed in an SBC of this type to reliably operate in a typical military computing environment, which will be discussed in the vehicle-mounted and airborne surveillance examples later.

Processor power and capability versus SBC form factor

The thermal characteristics of SBC electrical components vary greatly and drive the SBC form factor choice as well as the military computer casework's design. We are all familiar with Moore's Law about the doubling of the number of transistors that can be inexpensively placed on a piece of silicon every two years. One can argue the specifics of this trend and what the future holds, but one of the less-talked-about aspects of this "law" is the thermal issue related to adding more and more processing capability on a piece of silicon. Processors

Figure 1 | Small form factor and high-performance single board computers

with added capabilities such as two, four, six, and eight execution cores plus on-die memory cache sizes approaching 20 MB will generate more heat than previous generation CPUs. The maximum Thermal Design Power (TDP) and number of transistors per die shown in Table 1 illustrate how adding processor capability leads to increases in thermal design power.

Note the differences in the processor specs for the two general-purpose, long-life, embedded CPUs in Table 1. These two processors are used in very different military applications. The CPU in column 1 is used on a small form factor SBC deployed in an aircraft cockpit application, while column 2 illustrates a processor used in a dual-processor SHB design operating in rack-mount computer chassis on an airborne surveillance platform.

Also note the differences in TDP as it relates to the processor spec differences like the number of CPU cores and cache size. This difference in TDP is very typical and must be taken into account when designing a high-performance, dual-processor SBC and integrating the board design into a military computer.

The board designer needs to address these CPU design facts by:

1. Ensuring VRM power and ground traces contain adequate copper density.
2. Optimizing the SBC layout to maximize airflow over all board components to ensure adequate cooling at the specified operating temperatures.
3. Choosing or developing heat-sink designs that radiate the proper amount of heat away from the CPU die while maintaining the lowest stack-up height possible.
4. Placing SBC system memory DIMMs so they do not impede airflow.
5. Validating that the stated SBC operating temperature range has been qualified at the worst-case operating temperature ranges with a built-in safety margin of at least +5 °C.

	Intel Core2 Duo processor SU9300	Intel Xeon processor E5440
# of cores	2	4
# of threads	2	4
Clock speed	1.2 GHz	2.83 GHz
Cache	3 MB L2 cache	12 MB L2 cache
Lithography	45 nm	45 nm
Max TDP	10 W	80 W
Package size	22 mm x 22 mm	37.5 mm x 37.5 mm
Processing die size	107 mm^2	214 mm^2
# of processing die transistors	410 million	820 million
Sockets supported	BGA956	LGA771

Table 1 | Processor capabilities comparison (Source: Intel Corporation)

Let's take a look at two different military computer systems that utilize high-performance, dual-processor single board computers.

Vehicle-mounted system

Figure 2 illustrates how a dual-processor single board computer is integrated into a military computer designed for vehicle mounting.

The available vehicle mounting space, the intensive data communication requirements of the application, and the single board computer needed for the job drove form factor selection for this system, mounted in a Humvee.

The SBC is mounted with various hold-down brackets to withstand the shock and vibration of the application. This chassis has been tested and certified to the following MIL-STD-810F specifications:

- Mechanical shock – 40 *g*, 11 msec, saw tooth per MIL-STD-810F, Method 516.5
- Vibration – Operating – per MIL-STD-810F, Method 514.5, category 20
- Vibration – Nonoperating (transport) – per MIL-STD-810F, Method 514.5, cat. 5

Air filter selection is critical in this application because of the desert environments of the target application. Filter type and chassis construction were tested and certified via MIL-STD-810F, Method 510.4, Procedure I for dust-proof operation. Thus, this system illustrates one of the more rugged system platforms used with high-performance PICMG 1.3 SBCs.

Airborne system increases compute density

In airborne surveillance platforms, it is all about compute density. Typically SBCs with two multicore processors are desirable, and the more cores the better because the

Figure 2 | MIL-STD-810 validated system: A = Dual-processor SBC; B = Backplane; C = Plug-in cards; D = Drives; E = System connectors; F = Fan filters

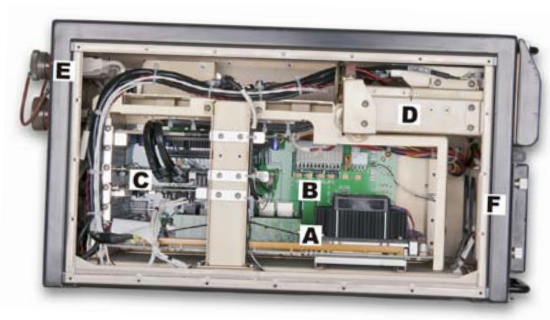

Figure 3 | Chassis with multiple high-performance SBCs

software applications take advantage of the large number of virtual machines made possible by each core. Further, the software applications also harness the ability of each core to process multi-threaded instructions and data sets. Placing two of these processors on a PICMG 1.3 SBC with four of these SBCs designed into a military rack-mount computer using a four-segment backplane results in 128 virtual machines in a single enclosure. Using these high-performance SBCs and clustering them in a rack-mount chassis as illustrated in Figure 3 saves on total payload weight by reducing the number of chassis enclosures inside the aircraft.

The airborne chassis shown in Figure 3 is made out of aluminum to further reduce system weight. Our Trenton Systems engineering team, in conjunction with the Tier 1 military OEM for this aircraft refresh program, developed redundant SBC hold-down brackets to ensure that the edge card fingers on the single board computers remain in place inside the backplane card slots during flight operations. This system illustrates how multiple PICMG 1.3 SBCs can be placed into a single computer enclosure to solve the problem of compute density.

The right form factor choice eases design challenges

There are a lot of single board computer form factors available for military and government computing applications. Sometimes a custom SBC design is required for the job and other times choosing a standard form factor from the plethora available today works just fine. The trick is to match the correct form factor to the application requirements, as far as compute density is concerned. Additionally, increasing processor capability generally means increases in heat generation. Fortunately, new processor die designs are mitigating this basic law of physics, but designers still need to do their system design homework to ensure a successful SBC implementation and to get the described SBC trends working in their favor. **MES**

__Brad Trent__ is Director of Engineering for Trenton Systems. Brad has held various development engineering positions in the embedded computing, industrial automation, computer workstation, and process system control industries. He holds a BS in Nuclear Engineering from The Georgia Institute of Technology in Atlanta. Contact him at btrent@TrentonSystems.com.

Trenton Systems | 770-287-3100 | www.TrentonSystems.com

SBCs move beyond traditional roles in embedded architectures

By Doug Patterson

Advances in IC density as well as the advent of new, low-power processors have enabled embedded designers to rethink the traditional role of the SBC within embedded military applications. With more functionality being housed on fewer components, valuable space is freed up on the board itself, allowing for deeper SBC integration within complete, fully functioning embedded computing subsystems.

U.S. Air Force photo by Staff Sgt. Joshua Strang

Time is money, and in embedded military applications under fire from budget cuts from sequestration, stalled programs, and aging technology, never has a statement been more true. Those system programs that thrive will be the ones that take advantage of the lower Size, Weight, and Power (SWaP) of today's embedded components and subsystems while successfully integrating network architectures at both the hardware and software levels.

One computing area in particular where this change is taking place is Single Board Computers (SBCs), especially in ground and airborne military systems. Manufacturers have begun to look beyond traditional SBC roles to see how embedded computing systems can benefit from the overall technological advancements within the industry. Instead of merely incorporating the latest iteration of a processing technology or of a smaller, denser IC, companies are taking a step back and looking to see what else can be done from an overall board, and even system, perspective.

Moore's Law is nothing new; circuits have been doubling in density and decreasing in size for years. The truly revolutionary part is that as the decrease in component size has allowed integrating several chips into just one, tremendous amounts of real estate have been freed up on the SBC.

And what is a designer to do with that extra space left on a board? Add more functions that couldn't previously be placed on the board, of course. This has enabled an SBC to move beyond its place as a mere component in a system that housed the processor and memory chips to an integral system component that is far more capable than its standard definition.

But the price to pay with all these added I/O features is that there are no backplane I/O pins to bring these functions off the SBCs. This means that new and updated industry mechanical and bus protocol standards are needed to keep up with the compression pressures from the silicon and board vendors.

The trend is clear: As more functionality is pulled onto SBCs, deeper integration at the subsystem level becomes more easily obtainable. The SBC can now handle many of the computing functions that could not be housed on the board before, including lots more on-board memory, I/O functions, interfaces, and control logic.

" ... Today's systems and platform integrators are quickly realizing that they are no longer tied to traditional embedded computing SBCs, backplane architectures, and system slots. "

In fact, today's systems and platform integrators are quickly realizing that they are no longer tied to traditional embedded computing SBCs, backplane architectures, and system slots; they can house all this functionality on one small subsystem that can be integrated onto one board, housed directly in a stand-alone, rugged small form factor enclosure.

Deeper integration of the SBC

SBCs are now moving further into the embedded computing integration process, becoming the central component in fully integrated subsystems. The board itself is now the keeper of far more system functionality than ever before.

In fact, IC density and package size have decreased so much that it is no longer the dimensions of the components that are driving the size of a military embedded system. Instead, overall size is now determined by the larger, rugged connectors required on the chassis/enclosure front panels. These larger connectors are needed to interface with the connector harnesses of many legacy platforms, which are currently under a technology-insertion, capabilities (or Spiral) upgrade program.

As the density of embedded components – and, therefore, SBCs – increases, there is a shift in how an SBC functions within an embedded computer. Smarter components equate to more intelligent systems that can now provide unique attributes to mission-critical computing, such as platform location monitoring built into the onboard FPGA of an I/O expansion unit, with a highly dense SBC at the core. This helps in terms of both form and function, since more functionality is contained within a smaller, compact overall system.

How low can you go?

With the recent technology and machine control precision advances in Direct Step on Wafer (DSW) hyper-Numerical Aperture (NA) immersion lithography now edging below 22 nm, continual die shrinks as seen in today's Very Large Scale Integration (VLSI) ICs are constantly tugging on Moore's Law (Figure 1).

This enhanced capability is benefiting SBCs by advancing the general availability of multi-GHz, multicore processors with dedicated integrated instruction and data caches, fat pipe fabric crossbar switches of high-speed serial lanes to high-speed memory buses, and on-chip real-world I/O.

These improvements to raw processor and memory access make the trade-offs of speed and performance less burdensome versus (instruction and data) clock cycles. It is now much easier to implement in software what was once rooted in the base silicon only a few years ago.

The new raw and enhanced performance allows SBCs to house non-generic control systems like digital signal processing, data security, and "on-the-fly" advanced "trusted platform" Data Encryption Standards (DES) and data cryptography. The trends and advances in this area that go above, and well beyond, the 128-bit portable encryption keys traditionally in use today are clearly seen on the horizon.

Balancing the benefits of COTS

Because they reduce overall design and integration time, COTS products continue to provide a cost-effective means of easily integrating advanced technologies in embedded systems, especially when implemented by manufacturers familiar with the nuances of military, defense, and aerospace applications.

Semiconductor Manufacturing Processes	
Line Geometries (nanometers, nm)	**~Year**
10,000 (10μm)	1971
3,000 (3μm)	1975
1,500 (1.5μm)	1982
1,000 (1μm)	1985
800	1989
600	1994
350	1995
250	1998
180	1999
130	2000
90	2002
65	2006
45	2008
32	2010
22	2011
16	2013*
11	2015*
6	2020*
4	2022*
* = projected	

Figure 1 | Projected line geometries in semiconductor manufacturing near 4 nm within the next 10 years.

In light of the state of many current military programs, COTS still provides a cost- and schedule-effective means to update technologies in existing programs, particularly in avionics and vetronics. However, power consumption and heat dissipation still remain large issues in these areas.

COTS-based SBCs primarily utilize multicore RISC and Graphics Processor Unit (GPU) processors from three main industry sources: Freescale, Intel, and AMD. These processors are surrounded by heaps of super-fast caches, fast DDR "X" SDRAM, serial fabric pipelines, and loads of platform-relevant I/O, all packed onto a single 3U or 6U Eurocard with daughterboard interconnect buses like VMEbus, CompactPCI, or OpenVPX and standard mezzanines.

The natural outcome – or side effect – of these larger SBCs with their 2+ GHz VLSI RISC and GPU processors is much higher power consumption, creating significant heat generation. It's not uncommon for these processors alone to dissipate well over 60 W. Add memory, I/O, and ancillary support circuitry for these boards, and dissipation can quickly reach over 100 W per backplane slot.

While an SBC can use new and clever heat removal methods, such as a copper-aluminum hybrid thermal frame, there's an associated cost and power density load put on the overall housing of these power-hungry processors and the subsystem boards they reside upon; this forces more costly liquid and direct impinged, spray-cooling solutions.

ARM lends a hand for military system designers

The best solution overall would be to minimize heat generation on these tightly packed SBCs from the start, eliminating the need for additional system cooling techniques. Encroaching quickly into the embedded computing space are ARM single- and multicore processors

These processors provide the equivalent raw processing horsepower as other more high-power processors currently used on military embedded SBCs, but with advanced and sophisticated on-chip and on-board power management; thus, they dissipate only 1/50 to 1/100 of the power of the larger VLSI RISC and GPU devices.

With these new ARM processors on next-generation SBCs and Data Concentrator Units (DCUs), imagine what a distributed and networked cluster of these remote interface/remote I/O engine nodes could do spread around tomorrow's military ground, air, or space-based manned and unmanned robotic vehicles.

Look closely at the latest smartphone to witness firsthand the power and capability of these ARM-based single core and multicore machines. Today's Android or iPhone contains enough raw processing performance to not only provide a full-featured cell phone, but also provide dual multimegapixel digital cameras and a high-definition video camera, a highly functional user touch interface, an HD display, and myriad other features – now too numerous to all be listed here.

Figure 2 | Units such as Aitech's A175 and NightHawk RCU are bridging the gap between ultra-high-performance SBCs and extremely low-power RIOs.

Along with all this, and while carrying 64 GB of memory, these processors:

- Execute all the individual video and audio codecs and modems needed for GPS, WAN, Bluetooth, and WiFi wireless;
- Conduct real-time monitoring of internal three-axis accelerometers and gyros, field intensity, and Hall effect sensors (for relative and absolute global positional and situational awareness);
- Directly manage all compression and decompression algorithms; and
- Operate an MP3 audio player, still and video camera, Web browser, and so on.

The list goes on, and this raw horsepower can run for several days on a battery pack measuring a volume of less than 0.5"^{3} and is barely warm to the touch, even under the heaviest usage and loading.

The military applications that will benefit from these low-power processors are as endless as the imagination of the engineer emboldened with the charter to use this new and exciting technology.

Integrated SBCs provide needed SWaP advantage

High-performance SBCs are stepping up to meet the challenges of today's and tomorrow's military and avionic Remote I/O (RIOs), DCUs, and Rugged Computer Units (RCUs) (Figure 2). These systems can marry ultra-high-power, super-performance SBCs with the super-low-power ARM-based RIOs, giving designers one fully functioning system.

And, the miniature scale of microelectronics today is displacing traditional architectures, enabling the SBC functionality to reside in one large FPGA, rather than require a separate board.

Take, for example, a networked, FPGA-based DCU used in highly data-centric environments with an internal SBC embedded within a large FPGA. Applications can include a remote interface data concentrator, engine and power train control, and data monitoring as well as vehicle prognostics data collection and Condition Based Maintenance (CBM).

FPGA-based DCUs are ideal for these types of applications when coupled with a large variety of I/O interfaces and the large user-programmable FPGA for the application Operational Flight Program (OFP). This "SBC embedded in an

FPGA" implementation provides a hardware and silicon, processor-specific technology isolation layer – essentially now processor agnostic and free from the processor component obsolescence that can cripple the long-term availability of a fielded subsystem.

Optimized for low SWaP, these small, lightweight DCUs and Remote Interface Units (RIUs) – with their self-contained "all-in-one" SBC – are highly useful in other environments: PC-based data concentrators and machine interface applications in manned and unmanned, ground, airborne, or space-based vehicles.

For military tracked and wheeled vehicle applications, SBC-based RIUs and DCUs can monitor key vehicle parameters and provide CBM functionality, thereby reducing the overhead costs of expensive, and most times unnecessary, periodic and preventative vehicle maintenance procedures.

When new, low-power, high-performance CMOS ARM and low-power Intel/AMD processors are paired with Power over Ethernet (PoE), these low-SWaP DCUs/RIUs can now be constructed and connected as flexible neural networks, meeting the system demands of tomorrow's defense vehicles and platforms with less power harnessing, fewer connectors, and, therefore, higher reliability.

The overall weight reduction – combined with a slimmer profile and natural convection/radiation cooling in units that can dissipate from less than 1 W to no more than 20 W – makes these units ideal for a variety of industrial and military and aerospace environments.

Reliable computing for today's critical environments

While the budgets of many military and avionics programs have come under fire, embedded computing manufacturers have stepped up to deliver integrated systems that help save both installation and maintenance time as well as costs. By combining higher-density components with low-power, cost-effective processing technologies, today's SBCs can provide the enhanced functionality and performance requirements demanded by critical defense systems. **MES**

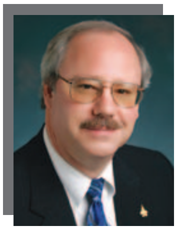

Doug Patterson *is VP, Military & Aerospace Business Sector, for Aitech Defense Systems, Inc. Doug has more than 20 years of experience in marketing, business development, and product management, including technical and communications accomplishments in telecommunications and harsh environment electronics. He also holds three patents in advanced metered mailing systems and nonvolatile memory redundancy mapping in severe/harsh environments. Doug holds a BSEE from BEI/Sacred Heart University. He can be contacted at doug.patterson@rugged.com.*

Aitech Defense Systems, Inc. | 888-248-3248 | www.rugged.com

Can sensitive data for tactical military environments be protected in the cloud?

By Robert Day and Louise Funke

When storing, accessing, and disseminating military data in the cloud, top concerns include security, data reliability and redundancy, and data location. The good news is that these can be delivered when secure virtualization pairs with a distributed cloud computing scenario.

U.S. Marine Corps photo by Sgt. Christopher R. Rye

While the promise of cloud computing, with its lower costs and improved access through utility computing and storage, is very attractive, it is currently difficult to achieve for users with highly sensitive data.

Consider the case of a warfighter who has identified a threat that must be reported. The warfighter sends the data directly back to his command where others can use it to take action and prevent casualties. Shortly after the data is sent, the warfighter is injured. A medic scans the warfighter's dog tag and accesses his medical history. The medic stabilizes him and inputs answers to several key medical questions that enable prompt, effective treatment after transport. While the warfighter is being treated, the data he sent earlier has been analyzed and tactical plans have been developed. On a carrier hundreds of miles away, pilots are being briefed about their target.

Key to the success of the aforementioned scenario is not only the way that military personnel access and input data but how easily that information can be stored and communicated to interested parties, without compromising the security of the information. Today's military is striving to enable warfighters with useful, actionable data and also equip them with the tools to capture and report critical data. In essence, there are efforts to push data access to the front, where data can be most effectively used. Data capture is also being enabled at the front, from the warfighter and from sensors, in order to obtain the most recent and useful data.

A natural way to further this approach is through some form of non-public cloud. A cloud approach – whether private, community, or a hybrid – would provide a host of benefits, including significant cost savings and increased agility for military organizations. Yet there are multiple challenges to deploying these kinds of tactical solutions today using current cloud technologies. However, a distributed computing approach to secure virtualization provides a viable solution to concerns surrounding data's security, reliability, and location within a cloud computing environment for the military.

Gathering storm: Security in the cloud

Security remains the greatest concern about using the cloud, even for private and community clouds. Questions being raised include:

- If all our key data is in the cloud, won't it be a more tempting, target-rich environment for hackers?
- With key data in the cloud, what happens if the cloud environment is impacted by a natural or manmade disaster?
- How can we take advantage of the cost savings of the cloud while still maintaining the separation needed between data classifications: unclassified, secret, and top secret?

The good news is that through a creative combination of highly secure virtualization and distributed computing, technologies are already available to address these concerns.

While all data may be "in the cloud," it doesn't mean it needs to be kept in one location, either physical or virtual. One way to lower the attack footprint of a private cloud is to use a distributed computing approach. With a distributed approach, multiple physical data centers make up the cloud and data is spread among the servers at various locations. Data isn't replicated on each server, but rather shards, or pieces of each database, are spread across the servers as designated by redundancy and location policies created by the administrator. Because the data is not all in one location, it's more difficult for an unauthorized person to acquire meaningful data. For example, a database of key targets might be sharded so that the ID of a target is on a server at site A, the location of the target is on a server at site B, and the people associated with a target are on a server at site C.

Because each shard of data is in multiple locations as defined by the redundancy policy, if a site experiences a catastrophic failure, no data will be lost and users will be able to access data from nodes at other sites. With a distributed data approach, even if a cloud data center is attacked and all data is lost at that location, the system knows where all the replicas of each shard of data are located and the system continues to operate without that data center. The system also recognizes that additional replicas of the shards that were stored at that data center must be created to adhere to the redundancy policy. As an example, the target data entered by the warfighter may have been stored in a nearby cloud server, or node. If that node was destroyed shortly thereafter, the target data would not be lost, as replicas were created and stored on multiple servers immediately after the data were entered.

While distributed computing improves security for cloud-based data, an extra-secure virtualization technology is required to fully realize the cost savings of cloud computing and the ability to host multiple networks on a single system. Secure software virtualization was created to address the needs of tactical military systems that require information and applications operating at different security levels to securely coexist on a single hardware platform. This removes the need for the costly deployment of multiple computer systems to facilitate communications and information from different forces or different intelligence levels in the battlefield.

Virtualization has become a major enabling technology for moving to the cloud by allowing multiple applications to co-reside on a single server platform and efficiently serve different types of data and applications to clients that connect to it. Size, Weight, Power, and Cost (SWaP-C) are usually improved with virtualized systems, which can be critical in field deployments. However, in a typical virtualized system, much of the virtualization of memory and devices is held in the same hypervisor code; hence, any breach of that code gives access to all of the memory and devices on that physical system.

Clearly this approach is not secure enough to allow different types or levels of sensitive information to reside on a single system. By using secure virtualization (see sidebar), true separation of memory and devices is key, and allows for different applications to securely coexist. In our example, the aircraft carrier has space constraints for computing systems and this favors the ability to have applications with multiple security classifications on the same system (Figure 1). The top-secret target information can be securely stored on the same system as the warfighter's medical data because they are completely compartmentalized.

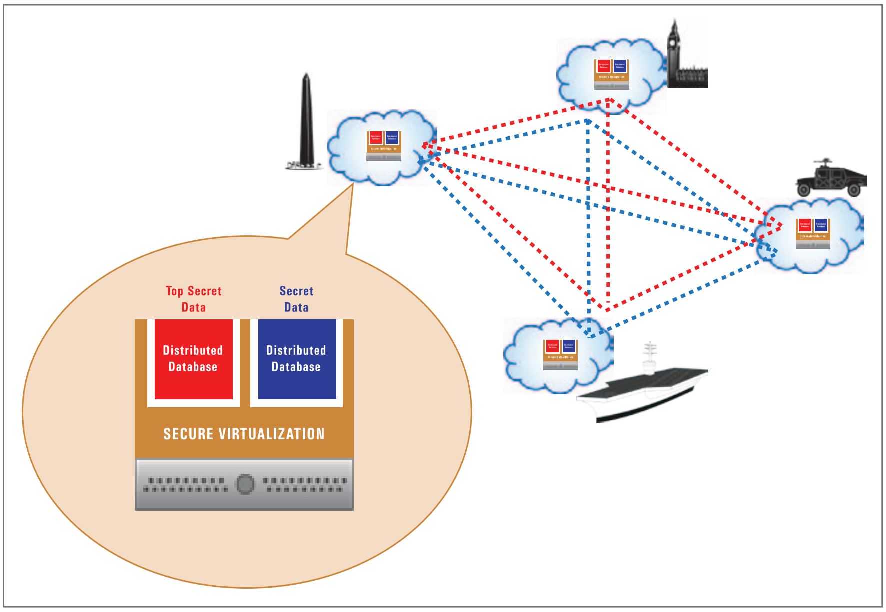

Figure 1 | Tactical military clouds can be used to store and share data at multiple classification levels using secure virtualization and distributed data management to maintain data security and resilience.

A closer look: Resilient data through distributed clouds

Right behind security, the next major concern for military organizations considering a move to the cloud is resilience. Cloud computing offers the potential of greater application and data availability yet also involves new risks as seen in the frequent outages in the commercial sector. Several multiday outages affecting large numbers of users have been well publicized in the past couple of years. Clearly, tactical military clouds need very high availability. In our example, the inability to get target information or warfighter medical data when needed would be unacceptable.

When it comes to system reliability, a distributed approach for applications and data enables high availability and resilience in the cloud. In place of the typical centralized application infrastructure stack – storage, relational database, application runtimes, and load balancing – these capabilities are instead cooperatively provided by identical nodes that may be deployed in a cloud, whether private or community.

With this design, the application infrastructure can be placed wherever needed around the globe, enabling mobile deployments, such as ships or vehicles, and locations near the front, to reap the benefits of the cloud. Scaling the system is as simple as adding nodes where they are needed.

Anatomy of secure virtualization

Anatomy of secure virtualization

The compartmentalization of data is achieved through two components of the system:

1. A separation kernel that securely partitions memory and peripheral devices that reside on the same physical hardware platform
2. Virtualization technology in the form of a secure hypervisor that provides multiple secure virtual environments

The combination of these two components allows multiple operating systems, applications, and data to reside in securely partitioned domains on the same hardware. By using a software separation kernel – which separates both applications and data at the lowest level and is linked to processor-based virtualization technology (for example, Intel VT-x and VT-d) – a single platform can host multiple secure virtual machines with no risk of compromise and mingling of applications, devices, or information.

Sidebar 1 | Secure virtualization has two primary must-haves: a separation kernel and secure hypervisor.

RUGGED COMPUTING SYSTEMS FOR OEM

"ZEUS" SZC91X

Rugged, Six-Way Independent Virtual Machine System, Westmere® CPU w/96GB RAM, Dual Removable Drives for up to 2TB of storage. Ultra-Small over all footprint only 13"x7.5"x4"

"GOLDEN-EYE III" S902R4

Rugged, Small, Lightweight, Quad Core Intel® i7 System w/16GB of RAM and Quad Removable SSD for up to 4TB of storage. Small 7"x5.5"x4" and 6.5lbs. (w/ Four Drives)
(Available in 2 drive configuration)

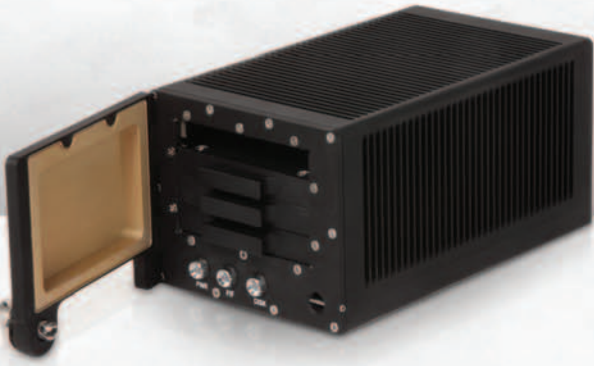

"DEPOT" SX401R4 NAS

Rugged, Ultra-Small, Low-Power Intel® Atom™ System w/2GB of RAM and Quad Sealed Removable SSD for up to 2TB of storage. Ultra-Small and Lightweight only 6"x3.5"x3" and 2.5 lb

"ARMOR" S822D

Rugged, Core 2 Duo® Penryn System w/8GB of RAM w/Ultra-Rugged Display and Dual Removable SSD up to 2TB of storage. Ultra-Small only 8.5"x13.25"x2.75"

"AXIOM" S279X

Rugged, Ultra-High Performance, Dual 3GHz Quad Core Xeon® System w/up to 32GB of RAM and Dual Removable SSD up to 2TB. Ultra-Small 11"x7.5"x1.75"

GMS COMPUTING ENGINES

General Micro Systems, Inc.
800-307-4863 **gms**4sbc.com
Since 1979

AS9100 REGISTERED FIRM

All Rugged products are fully-sealed and Extended temp -40°C to +85°C Compliant to MIL-STD-810G, MIL-STD-704E and MIL-STD-461F

All GMS products are proudly designed and manufactured in the U.S.A. and are fully ITAR compliant. © 2011, General Micro Systems, Inc.

The "brain" of such a system, the geographically distributed relational database, resides in multiple locations and provides like-local application performance, yet functions as one entity. The database that contains the warfighter's medical history could benefit from this approach. The information appears and acts as though it is being processed from a centralized database even though it is geographically distributed. The processing is real time and satisfies Atomicity, Consistency, Isolation, Durability (ACID) guarantees. There is no active-passive concept for the nodes; all nodes are equal, and there is no "master node."

Because the data is stored redundantly across nodes based on policy, the system can continue processing if a node or site fails, while automatically rebuilding redundancy. The self-healing architecture recognizes possible problems and automatically adjusts to prevent a disruption in service. For instance, if the node usually accessed by the warfighter failed, he would be automatically redirected to another node in the system.

Moreover, because all nodes provide the required application services in a resilient fashion, organizations no longer need to set up and maintain dedicated failover sites. No resources are dedicated purely to disaster recovery; instead, surplus resources are used to satisfy increases in application demand and improve performance for application users. When resources are deployed redundantly, failures don't result in outages and five-nines availability is achieved. The result is that a well designed military cloud can be incredibly resilient and highly available.

Location, location, location for cloud data

Another hurdle for military organizations considering cloud solutions is the currently limited ability to control the location of data in the cloud. Some DoD organizations might require the ability to restrict the data's location. For instance, data shared by coalition forces may be of the same classification, yet each of the coalition partners may have data that should not exit their facilities.

Data governance must be an integral part of a distributed database and application system in order to reduce the cost and risk associated with compliance management. Data location policies allow administrators to establish rules for where data can or cannot be stored.

For example, by establishing location policy rules, administrators can specify that certain tables or portions of tables must or must not be stored in various locations. If critical data must be stored only in a particular region, or may not be stored at locations with inferior physical security, administrators can pinpoint these restrictions. Policies can also be designed so that data rules are enforced as data is created and updated, ensuring that systems are always in compliance with established policies.

Deployment options

If both private and community clouds are planned, data may be initially deployed on private clouds. As community members become comfortable with the security and policy measures, data can be moved to community clouds where it can be shared. An even more conservative approach would be to begin with a mix of on-premises nodes and private cloud nodes, migrating data to private clouds as confidence in the cloud infrastructure grows, and finally to community clouds as appropriate.

Secure virtualization enables cloud-based systems

When the resilience and data jurisdiction offered by distributed data technology are coupled with secure virtualization, the true economies of scale of a cloud-based approach can be realized. Virtualization is critical to successful cloud deployments, especially those in field-hosted environments, because of SWaP-C reduction for each deployment. However, secure virtualization is needed to enable cloud-based systems to handle sensitive data. The additional compartmentalization provided by a secure virtualization platform allows data processing and storage at different security classifications on a single hardware platform.

Secure virtualization along with data location control and very high availability are key to mission-critical military deployments that support multiple data classifications. Cloud solutions designed to protect sensitive data for these kinds of tactical military environments are now attainable using secure virtualization in a distributed computing environment. **MES**

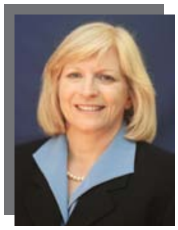

***Louise Funke** is VP of Marketing at TransLattice and has extensive experience in the marketing and delivery of technology solutions. Prior to TransLattice, she served as VP of Marketing for BlackHog, a supply chain software developer. Previously, at Sun Microsystems, Louise managed product introduction teams for the federal market. She holds an MBA from the Stanford Graduate School of Business and a BS in Industrial Engineering from U.C. Berkeley. She can be contacted at lfunke@translattice.com.*

TransLattice
408-749-8478
www.translattice.com

***Robert Day** is Vice President of Marketing for LynuxWorks. With more than 20 years in the embedded industry, his most recent position prior to joining LynuxWorks was heading marketing for Mentor Graphics' embedded software division. Based in San José, California, Robert is a graduate of The University of Brighton, England, where he earned a Bachelor of Science degree in Computer Science. He can be contacted at rday@lnxw.com.*

LynuxWorks
408-979-3900
www.lnxw.com

U.S. Air Force photo by Tech. Sgt. Sean M. Worrell

Secure virtualization for tactical environments

By David Egts

Virtualization has proven its value in the data center, but can it work in tactical environments? Yes, but only if it's secure.

The evolution of computing capabilities to allow the mobilization of data has paved the way for the transformation of tactical vehicles into rolling data centers. These impressive machines have provided invaluable resources for warfighters to dominate the battlefield. Unfortunately, despite the obvious benefits to packing a Humvee or tank with real-time intelligence computing equipment, these capabilities come at a not-so-obvious price: a sacrifice of precious power, cooling capabilities, and logistical space to accommodate the new equipment. For example, too many computers could adversely impact visibility or generate additional heat in an already scorching environment, dulling a soldiers' reaction time and potentially increasing health and safety risks.

Interestingly, these rolling data centers are experiencing many of the same challenges found in today's enterprise data centers – again, power drain, insufficient cooling, and lack of space. To address these challenges, data center managers have turned to virtualization to perform physical to virtual server consolidation, regularly achieving a 10 to 1 optimization or better. Can the lessons learned by virtualization in the data center also benefit tactical vehicle environments? What about special security requirements of tactical vehicles that contain content of differing classification levels? It is essential to understand the unique challenges and similarities between virtualized and physical server consolidation, as well as some of the secure virtualization advances made by the open source community that can be directly applied to help the warfighter.

Benefits of virtualization

There are several benefits to the deployment of virtualization to the field and in data centers, including consolidation, uniformity, live migration, and performance, which can specifically aid warfighters and increase efficiency on the front lines.

Consolidation

When workloads are consolidated with virtualization, hardware utilization increases. More can be done with less hardware, resulting in power, space, and cooling savings. In tactical environments, capabilities can be added with nominal increases in compute requirements. By consolidating workloads on fewer physical systems, the newly available space can be repurposed for additional mission features, more bullets, and literally more elbow room for the crew.

Uniformity

Virtualization allows data center managers to abstract hardware from the OS. This allows the data center manager to acquire the best performing hardware, for the lowest possible price, without needing to recertify software stacks. Hardware vendors can be pitted against one another to deliver the best value. In tactical environments, this advantage is even more profound, with no need for each application to have its own computer with an exotic form factor, power requirements, proprietary connectors, and so on. As a result, new capabilities can be added much more easily and hardware refreshes can happen much more quickly because, as mentioned, recertification time is reduced. Additionally, not every vehicle may need the same capability, so the need to rigidly budget power, space, and cooling, whether the capability is used or not, becomes relaxed. In the end, uniformity makes the warfighter much more nimble.

Live migration

If imminent hardware failure is detected on a hypervisor in the data center, a workload can be live migrated with zero downtime. Even if the hypervisor fails unexpectedly, the virtual machines running on it can be restarted on other hardware without user intervention. In tactical environments, battle damage resulting in the loss of a computer may result in loss of critical capabilities needed in the heat of battle. By running workloads directly on physical hardware, the capability may not be restored until the warfighter returns for the vehicle to be serviced, which may be too late. If the workload is virtualized, it could be migrated in real-time to or restarted on a functional server elsewhere in the vehicle without warfighter intervention. As a result, the probability of the warfighter losing a mission-critical component during battle for an extended period of time is lessened.

Performance

Some workloads that require low-latency or high-bandwidth communication between systems actually run faster when virtualized on the same piece of hardware. This is because the virtual network communication is done over the system backplane instead of going from physical system to physical system across a much slower and higher latency networking infrastructure. This increased performance and lower latency can help the warfighter respond much more quickly in the heat of battle.

Is virtualization secure?

Not by itself. If a hypervisor is compromised, not only is it compromised, but so are the virtual machines running on it as well as the virtual machine disk files connected to the hypervisor that may be running on it, on another hypervisor, or at rest. When a hypervisor exploit happens, restoring all virtual disk files from a known reliable backup is needed. If the time of exploit cannot be pinpointed, a full reinstall from scratch may be necessary. Many virtualization security solutions put in place offer a trusted means for the virtual machines to communicate with each other, but these tools offer no protection in the event of a hypervisor compromise. If virtualization is to be deployed in tactical environments, the hypervisors must be secure so the warfighter can trust the integrity of the components supporting the mission. So do technologies exist to prevent this type of exploit? Yes, thanks to the open source community.

Secure virtualization with sVirt: *sVirt* is a secure virtualization open source project, built upon a time-tested SELinux open source project pioneered by the NSA and the commercial software industry, dating back to 1999. sVirt is used to securely isolate KVMs (or Kernel-based Virtual Machines) from the hypervisor and from each other. Before we get into the details of sVirt, let's do a quick review of SELinux and KVM:

SELinux predates x86-based virtualization and was originally used to lock down physical systems to prevent rogue applications from harming one another as well as the main system. SELinux is all about labeling. Every object on an SELinux system is labeled according to its function. This includes processes, users, files, hardware devices, network ports and adapters, and so on. The SELinux policy allows objects with certain labels to explicitly access objects with other labels and deny everything else. For instance, if a rogue web server process wants to read a password file, SELinux will block and log the access since the SELinux policy does not explicitly allow processes with a web server label to access files labeled as password files. SELinux is not an add-on to Linux – it is built in. SELinux has been enabled by default in mainstream enterprise Linux distributions since 2007.

KVM is an open source project that emerged in 2006 and was accepted in the upstream Linux kernel in 2007 and can today be found in numerous

Linux distributions. KVM is unique in that all KVM virtual machines are Linux processes. As such, KVM is very lean because it can leverage the time-tested Linux operating system to manage VMs, perform hardware enablement, and leverage advances in power efficiency and performance, to name a few. As a tremendous side effect, in the same way a web server process can be confined with SELinux policy, a KVM virtual machine process with the SELinux policy called *sVirt,* mentioned previously, can also be confined.

At a high level, sVirt works by placing SELinux "force fields" around virtual machines, which confine what they can do. sVirt can even protect virtual machine guests that do not have the built-in protections of SELinux like Windows, as pictured in Figure 1.

When a virtual guest starts, the KVM Linux process is given a unique SELinux label by sVirt. Correspondingly, the virtual disk file of that virtual machine is given a matching label. With the exception of the process and the virtual disk file, no other object on the system has the same label. When a second virtual machine starts, the second virtual machine's KVM Linux process and virtual disk file are given their own unique but matching labels. This is all pictured in Figure 2.

In the event that a virtual machine is able to do a hypervisor exploit and gain administrator access, sVirt at the kernel level confines the virtual machine to only what is allowed by the sVirt SELinux policy, namely to see its own disk file (not others), use other resources narrowly allowed by virtual machines, and nothing else.

Sounds great, but does sVirt really work? Actually, yes. At the Black Hat USA 2011 conference, Nelson Elhage presented a method to compromise KVM, but the use of sVirt defeated his approach.

Side note: Interestingly enough, the labels created by sVirt are Multi Category Security (MCS) labels. MCS (and Multi Level Security or MLS) labels are concepts derived directly from secure computing efforts driven by the NSA and other agencies. The fascinating part is that this demonstrates how agencies with classified data requirements have helped the open source community take their requirements from niche to mainstream. By leveraging the proven MCS capabilities of SELinux, Common Criteria certification of KVM with sVirt becomes much easier.

Red Hat is a contributor on both the sVirt and KVM projects, as well as SELinux, utilizing their Red Hat Enterprise Linux and Red Hat Enterprise Virtualization products. For more information on these projects, visit www.selinuxproject.org or www.linux-kvm.org.

Future work

Things can always be made better and more secure. An additional concept being integrated into virtualization is the concept of resource control groups or *cgroups*. In the same way as SELinux can confine access by Linux processes, cgroups is used to control the CPU, network, memory, and other resources consumed by Linux processes. Since

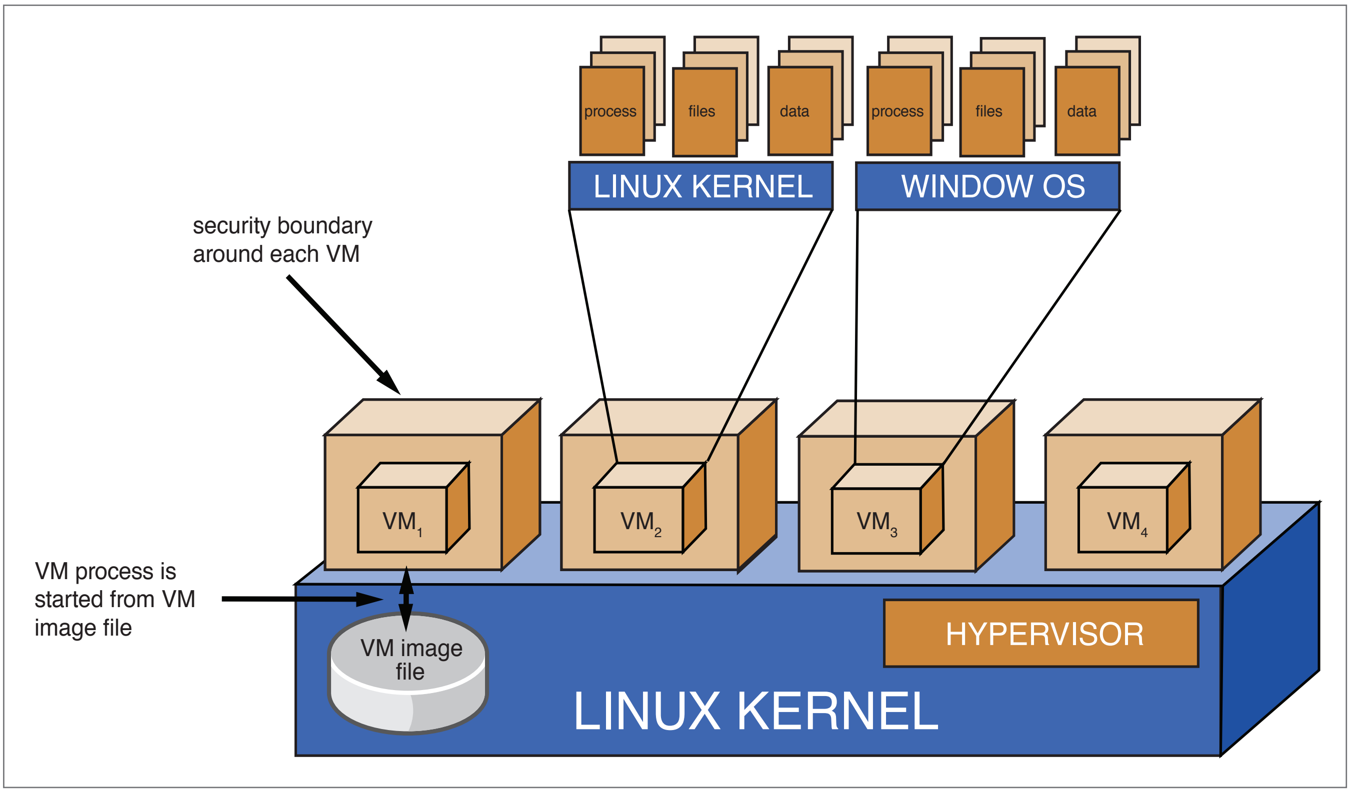

Figure 1 | High-level overview of KVM with sVirt protection

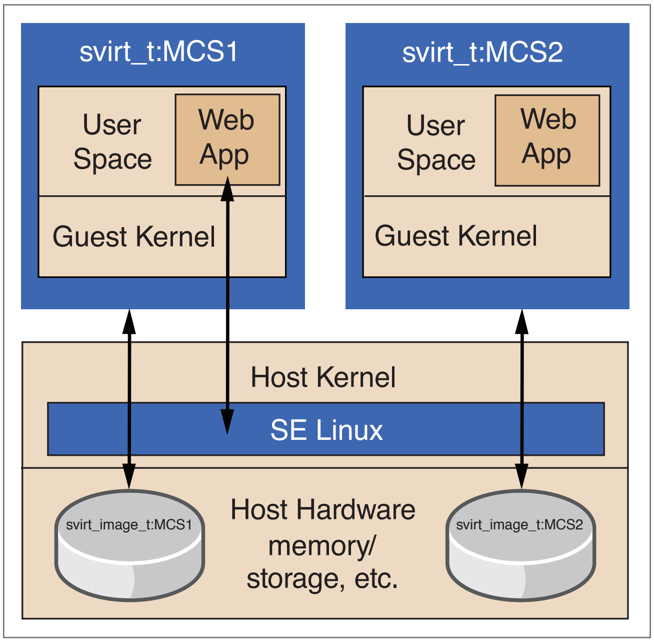

Figure 2 | Unique virtual machine process and disk file dynamic labeling by sVirt

KVM virtual machines are Linux processes, they can also be confined by cgroups in the same way they are confined by sVirt. This paves the way for resource fairness when it comes to virtualization multitenancy. One integrator's virtual machine won't overconsume their fair share of system resources, which could adversely affect another integrator's virtual machine on the same hypervisor. From a security perspective, cgroups can also be used to prevent a denial of service attack on one virtual machine from taking down an entire hypervisor.

Another open source effort under consideration is that of trusted computing. As computing resources containing sensitive data are pushed deeper and deeper into the battlefield, a secure means to boot and trust federated systems that aren't locked and under guard in a data center becomes paramount. The open source community along with government agencies and hardware and software vendors are actively working on trusted computing technologies to address these needs. Since KVMs are Linux processes, the ability to rapidly adopt these emerging trusted computing methodologies becomes much more straightforward as opposed to inventing something wholly new. **MES**

David Egts is a Principal Architect at Red Hat, Inc., specializing in the application of open source enterprise infrastructure technologies within federal, state, and local government agencies, the Department of Defense, and educational institutions. Contact Dave at degts@redhat.com.

Red Hat
703-748-2201
www.redhat.com

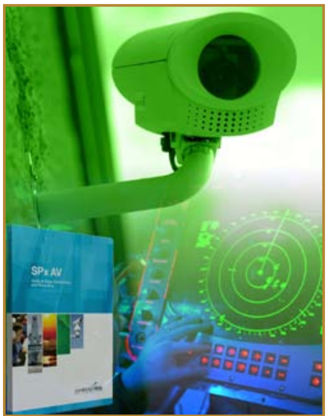

Radar software adds audio and video to its repertoire

Some say you can't get enough of a good thing, but in a slightly different vein, Cambridge Pixel showed us that you can improve an already-good thing. The company has added audio and video capability to its SPx AV software, a toolkit of components useful for radar in command and control, security, UAV, and airborne radar applications. With the updated product retaining its former nomenclature, the SPx AV software library now features video and audio distribution, capture, and recording savvy. The software library affords video and audio quality of service control, network streaming, compression, and real-time video display providing underlay, overlay, and watermarks.

Suited to radar video distribution, radar visualization, target tracking, and plot extraction, the SPx AV software additionally supports H.264 video compression on GPUs, to proffer the real-time display of several layers of video and graphics.

Meanwhile, the software utilizes an open-systems approach but allows user customization at the same time. And users of the former version of the software don't need to ditch what they've got: The new version can be used alongside it for multifunction display and sensor distribution technology utilization.

Cambridge Pixel | www.mil-embedded.com/p365945 | www.cambridgepixel.com

VPX/VITA 62 power supply enlivens mil systems

While the burgeoning VPX (VITA 46) VITA-spawned form factor is taking the military embedded realm by storm, an optimal power supply ensures optimal performance and reliability. Enter the VITA 62 Power Supply standard, which defines a power supply geared for VPX chassis while meeting rugged VITA 48 specifications. The good news is that vendors including Extreme Engineering Solutions are already producing VITA 62 power supplies, such as the company's XPm2120, ready to rev up those VPX wares.

The 3U VPX form factor XPm2120 power supply features an across-the-backplane primary distribution rail of 12 V to the payload modules. Meeting MIL-STD-704 for 28 VDC input power, the power supply proffers as much as 300 W running on 3.3, 5, or ±12 V with an efficiency rate up to 90 percent. EMI filtering in lockstep with MIL-STD-461E is also provided, as is a conduction-cooled operating temp of -40 °C to +85 °C at the thermal interface.

The .8" pitch, 1.1 lb power supply additionally features an IPMI controller to monitor temperatures and voltages of the board and to control output power. The controller is connected to a backplane, thanks to an I^2C system management bus.

Extreme Engineering Solutions (X-ES) | www.mil-embedded.com/p365309 | www.xes-inc.com

Rugged tablet prevents soldiers' mission knowledge draught

To keep soldiers on the ground in-the-know when it comes to mission strategy, wouldn't it be nice if they could view mission-critical data out in the field? A rugged tablet is certainly viable, and Panasonic Solutions Company's Toughbook H2 could be just the ticket. The fully rugged Windows-based tablet PC runs on the Intel Core i5-2557M vPro processor, speeding along at 1.7 GHz or up to 2.7 GHz on Turbo Boost. Standard on the device are a 320 GB, 7,200 rpm shock-mounted hard disk drive and 4 GB RAM, which can be increased to 8 GB. A solid-state drive offering 128 GB storage is optionally available.

With a battery life of 6.5 hours and a starting weight of 3.5 lbs, the Toughbook H2 features two identical hot-swappable batteries touted to charge rapidly. And the soldier in the field can surely benefit from the 10.1" dual-touch display, which is sunlight readable and renders touch-screen and digitizer functionality. Operating at temperatures of -20 °F to +140 °F (-28.9 °C to +60 °C), the tablet also withstands a 6-foot drop in accordance with MIL-STD-810G and has a sealed, all-weather design in a polycarbonate-encased magnesium alloy chassis. Connectivity is via Bluetooth or Wi-Fi 802.11, with optional 4G mobile broadband and optional Gobi 2000 3G mobile broadband. Meanwhile, interfaces include USB, Ethernet, and serial, while optional security features comprise insertable or contactless SmartCard readers, in addition to a fingerprint reader.

Panasonic Solutions Company | www.mil-embedded.com/p367718 | www.panasonic.com

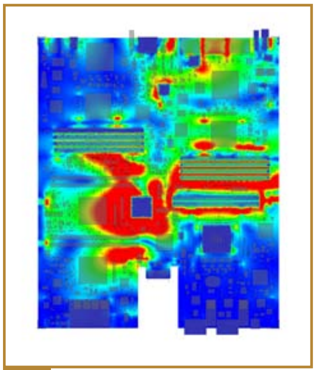

Software implements Physics of Failure Analysis

Physics of Failure Analysis is a topic that is critical to rapidly and cost-effectively delivering safe, highly reliable systems into military customers' hands. One company focusing on this paradigm is DfR Solutions, apparently, with its Sherlock Automated Design Analysis Tool software. Sherlock uses Physics of Failure Analysis to help designers spot the possible points of failure early in the design process, thereby eliminating the innate inefficiency of the build-test-fix cycle.

Sherlock works in three phases: 1) Data input – the same files that the user company would send to a manufacturer for board assembly can be used after data is easily imported into Sherlock. 2) Analysis – A holistic analysis is rendered, taking into account things such as time to failure, part failure, shock and vibe, and thermal, specifically: virtual thermal cycling, solder joint fatigue, virtual shock, Conductive Anodic Filaments (CAFs), virtual vibration, and Plated Through Hole (PTH) fatigue. 3) Report and recommend – Results are provided in several formats such as overlays, life curves, histograms, or in a tabular output. Reports are rendered in only a few minutes.

DfR Solutions | www.mil-embedded.com/p365589 | www.dfrsolutions.com

Recaster bridges secure enterprise and military LANs

The ability to transmit video content securely is a must for the military. Now VITEC's new release of its Optibase Ocaster MPEG recaster/reflector can do just that, acting as a LAN-to-LAN reflective bridge for video streaming between military LANs and enterprise LANs. The new Ocaster release features "full military encryption," the company says, including AES-128 support in addition to 256-bit encryption. Also included are Key-Length-Velue (KLV) and Cursor-on-Target (CoT) metadata in asynchronous and synchronous forms.

Ocaster works by converting several video streams into unicast streams, and vice versa, to ensure efficiency of network transmission. Closed captioning can also be used without hindering performance or speed of transmission. Ocaster's centralized, secure, Web-based Graphical User Interface (GUI) affords management of up to 50 MPEG video plus ad hoc setup and teardown. And finally, just about any IP stream can be processed in real time, as Ocaster is compatible with the MPEG-4 H.264 and MPEG-2 standards, and high-definition IPTV encoders. In any case, "an advanced traffic shaping algorithm ensures Quality of Service (QoS) for reflected/recast video payloads on any network and maintains consistent bandwidth utilization," the company adds.

Optibase, a VITEC Company | www.mil-embedded.com/p364914 | www.vitecmm.com

VPX goes HPEC

Some military programs need high-end processing, and that's all there is to it, for example, an airborne classified data system. Case in point: The Kontron-bred High Performance Embedded Computing (HPEC) platform, in an 8U VPX form factor, was selected for just such a system. And the reasons this ware stands out to us, too, are plenteous. Described as a "supercomputer-like system" by the company, the DSP-suited HPEC platform comprises 18 of Kontron's VX6060 Core i7 2 GHz (or more) computing nodes including 8 GB DDR3 memory on each board, in addition to 36 tightly coupled processors within the HPEC platform. Several switched fabric interconnects are housed in the backplane. And the 19" footprint HPEC platform – suited to radar, SIGINT, sonar, and video processing needs of UAV or aircraft programs – delivers 1.44 Teraflops, thwarting the old paradigm of military programs sacrificing performance (or facing a proprietary technology path).

Digging a bit deeper, the COTS-based HPEC platform utilizes a 10 GbE/PCI Express combo for the data plane, and the control plane is good old GbE; this setup yields 540 Gbps aggregate system throughput. Other components include PCIe switches, two 10 GbE mezzanine cards, PMC/XMC carrier boards, and GbE switch cards. Temperature ranges supported are conduction-cooled at -40 °C to +85 °C or air-cooled at 0 °C to +55 °C.

Kontron | www.mil-embedded.com/p367717 | www.kontron.com

Critical Recording in Any Arena

When You Can't Afford to Miss a Beat!

Introducing Pentek's expanded line of Talon™ 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 2.2 GB/sec
- Recording and playback operation
- Analog signal bandwidths to 1 GHz
- Shock and vibration resistant Solid State Drives
- GPS time and position stamping
- Hot-swappable storage to Windows® NTFS RAIDs
- Remote operation & multi-system synchronization
- SystemFlow® 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/mestalon for your FREE online *High-Speed Recording Systems Handbook* and *Talon Recording Systems Catalog*.

TALON

PENTEK

Setting the Standard for Digital Signal Processing

25 YEARS

Pentek, Inc., One Park Way, Upper Saddle River, NJ 07458 • Phone: 201.818.5900 • Fax: 201.818.5904 • e-mail:info@pentek.com • www.pentek.com