Military
EMBEDDED SYSTEMS

VOLUME 7 NUMBER 6
SEPTEMBER 2011

MIL-EMBEDDED.COM

INCLUDING:

John McHale
Unmanned applications still hot amidst looming defense cuts

Field Intelligence
Many-core heads for advanced sensors

Mil Tech Insider
COTS moves up integration food chain

ANNUAL RESOURCE GUIDE

Also:
Tactical communications:
Smartphone technology
for the battlefield

NASA turns to the open source community for advice

*An interview with Ray O'Brien,
CTO for IT at NASA Ames Research Center*

www.opensystemsmedia.com

SUBSCRIBE NOW

mil-embedded.com/subscribe

Military
EMBEDDED SYSTEMS

September 2011 Volume 7 Number 6

COLUMNS

DEPARTMENTS

RESOURCE GUIDE

Military EMBEDDED SYSTEMS

OpenSystems media.

Unmanned applications are still a hot market – even amidst pressure of looming defense cuts

By John McHale, Editorial Director

Hello, my name is John McHale and I'm the new editorial director for *Military Embedded Systems,* replacing founding editor Chris Ciufo, who has moved onto other endeavors in the embedded community.

Many of you might know me from my 15 years at *Military & Aerospace Electronics* magazine, where I not only covered embedded electronics, but also commercial avionics and homeland security technology.

The first day, week, or month on a new job can be exciting and stressful all at once, but my anxiety was quickly assuaged when I arrived at the Unmanned Systems North America show during my second week on the job and saw so many old friends. I was particularly heartened that so many of the exhibitors had positive things to say about *Military Embedded Systems* – the people and the product.

I guess I made the right choice. Whew!

When I first started in the industry, "embedded systems" was not a major buzzword in the military electronics community. At the time, many in the defense industry were still trying to puzzle out the meaning of the Commercial Off-the-Shelf (COTS) procurement mandate and how it would affect designs at the component and system levels.

Exhibitors at the first Embedded Systems Conference I attended in Chicago back in 1998 did not even have military applications at the top of their target list. However, I do remember engineers from a couple of COTS suppliers sharing a booth at the show, saying embedded technology was a strong growth area in the long term for military suppliers as future military system requirements would call for smaller sizes coupled with exponential increases in performance.

Incidentally, those two companies – Dy4 Systems and Virtual Prototypes – are still major suppliers of embedded technology for military systems. Today, they are better known as Curtiss-Wright Controls Embedded Computing and Presagis, respectively.

Both companies changed hands a few times over the past 15 years, which was not uncommon among COTS suppliers. However, what is so refreshing about this industry is that while company names change, many of the people remain or end up somewhere else in the industry. Remaining in the same industry when transitioning jobs is always a positive, but it helps when the market historically is healthy without the cyclical, crash and burn characteristics of markets such as telecom.

Those positive characteristics of long-term market health were reinforced to me at the Unmanned Systems event. Unmanned vehicles – be they Unmanned Aerial Vehicles (UAVs), Unmanned Ground Vehicles (UGVs), or Unmanned Undersea Vehicles (UUVs) – is one of the fastest-growing areas of defense spending.

All of these platforms require sophisticated electronics payloads to carry out their missions, whether they are for surveillance and reconnaissance or precision strikes. The payloads must fit into unique form factors and small packages for the many different-sized unmanned platforms; this means increasing opportunities for embedded designers of FPGA boards, digital signal processing systems, and RTOSs, to name a few.

Unmanned applications bode well for the long-term growth of the embedded systems market, but all military designers are at the mercy of Department of Defense spending. And with major cuts looming on the horizon, quite a few are nervous about the immediate future.

Exhibitors and attendees at the show were not calling it a doomsday scenario; however, if big cuts are made, they feel that it will seriously inhibit growth for the next few years. Research and development funding for new programs is at risk, which could result in short-term layoffs and even more industry consolidation.

Some AUVSI attendees were still positive despite budget-cut rumors, having signed long-term contracts and expressing confidence that no matter the money issues in Congress, the U.S. will need to find a way to maintain its status as the world's toughest military force. Global terrorism hasn't gone away, and China's military technology gets more impressive every day.

They also feel that major cuts could result in opportunities for retrofits maintenance of older systems.

As of this writing, though, it is still anyone's guess as to what will happen with the DoD budget. We will have more coverage on how COTS suppliers might be affected by the next budget in our October issue.

In the pages of this edition, we have our regular columns; an interview with Ray O'Brien of the NASA Ames Research Center on how NASA is seeking advice from the open source community; and our annual Resource Guide, which contains a plethora of profiles from military embedded suppliers with briefs on their latest products.

I'm looking forward to seeing many of you at upcoming industry events. Please feel free to contact me via email as well at jmchale@opensystemsmedia.com.

Lastly, I would like to thank Chris Ciufo for all his efforts in creating *Military Embedded Systems* magazine and in making it a regular read across the industry. We wish him well.

Thanks,
John McHale

Field Intelligence

Many-core processor devices head for advanced sensors

By Duncan Young

Advanced sensor processing continues to demand greater performance and increased frame rates, while consuming less of a vehicle's Size, Weight, and Power (SWaP) budget. The FPGA, and, more recently, the General Purpose Graphics Processor Unit (GPGPU) have been adopted for the highly parallel, repetitive front-end processing of raw sensor data. Both these device types use large arrays of simplified execution cores to continually process incoming sensor data streams. This many-core approach to enhanced processor performance is growing rapidly as new devices, such as Tilera's TILEPro64, become available. Driven by similarly data-intensive applications such as video transcoding, encryption, and deep packet inspection, the many-core tile offers another step up in performance for new sensor systems.

Transistor count

Silicon manufacturing processes are currently achieving upwards of 2 billion (2×10^9) transistors on a chip. However, as clock speeds are not increasing at the same rate, chip designers use the extra transistors to offer more functionality to increase their products' performance. Because of their legacy requirements, mainstream processors such as Intel's Core i7 tend to be improved incrementally by increasing the processing core count – by adding new graphics capability and introducing enhanced 256-bit vector processing capability for each core. Power dissipation becomes an issue as transistor count and clock speeds increase. Power must be managed and traded for ultimate performance, specifically in embedded applications where space and cooling capability are limited.

FPGAs have also benefited from the same process improvements, offering many hundreds of arithmetic cores and more on-chip memory and interconnects. However, an FPGA-based solution can be difficult to implement, requiring intimate device knowledge, but is ideal where

ultimate processing density is needed. An FPGA requires careful thermal and timing analysis, plus neither partial or online reprogramming have yet matured into realistic options for deployed capability growth. Despite the real advantages that FPGAs offer, many applications are turning toward software programmable alternatives for their rapid development and flexibility. The GPGPU and many-core tile processor demonstrate how fresh approaches can use the abundance of available transistors to achieve large arrays of cores for the parallel, multithreaded, repetitive operations needed at the front end of a sensor's processing chain.
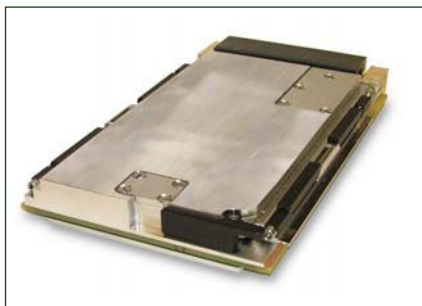
Figure 1 | MCP500 many-core single board computer from GE Intelligent Platforms

Data flow and processing strategies

Both GPGPU and many-core tiles are designed to process multiple data streams at very high rates. Generally, a GPGPU operates best on large, structured data sets running multiple threads of similar algorithms. By comparison, a tile device such as the TILEPro64 has 64 general-purpose processing cores, each with its own L1 and L2 caches, organized in an 8 x 8 matrix with nonblocking, high-speed switched data paths. This makes it well suited to content-based decision making and more varied use of the available cores. Supported by Linux, cores can run individual copies of the operating system, or many cores can be configured to share in a Symmetrical Multiprocessing (SMP) environment – or combinations of each are possible on the same device.

The tile's ability to use individual cores or to share data and run multiple threads in many cores makes it very efficient at raw radar data processing tasks: reorganization, beam forming, and adaptive cancellation. Today, such tile processors do not offer hardware floating-point capability. Software technologies can be adopted for less-demanding applications, although greater performance will require a heterogeneous architecture using a further processing stage based on GPGPUs or vector processors such as AltiVec or Intel's AVX, while still achieving overall savings in board count.

Video processing

Similar to the FPGA and GPGPU, the many-core tile processor has much greater military application potential than just radar. Sensor fusion and 360° local situational awareness in small agile platforms such as moving armored vehicles and helicopters are additional examples, requiring many channels of video data to be processed at high frame rates. Designed for many of these rugged, complex sensor processing applications, the MCP500 is a TILEPro64-based SBC from GE Intelligent Platforms, offered in 3U VPX and VPX-REDI form factors and providing 2 GB of DDR2 SDRAM, dual 10 Gbps XAUI Ethernet ports, and dual four-lane PCI Express interfaces (Figure 1).

Continuous process technology improvement still supports Moore's Law rates of growth in transistors, though legacy architecture and packaging technology can constrain mainstream processors from taking full advantage. However, the video, gaming, Internet, communications, financial, and entertainment markets all have the potential to spawn innovative many-core architectures providing additional scope in the future for sensor design improvement and diversity.

To learn more, e-mail Duncan at duncan_young1@sky.com.

Mil Tech Insider

PCOTS versus LRU: COTS moving up the integration food chain

By Steve Edwards

The COTS proposition, taking state-of-the-art commercial semiconductor technology and deploying it on open-standard military board architectures, is rapidly migrating up the technology food chain to the preconfigured subsystem level. Packaged COTS (PCOTS) subsystems are being driven by customer interest in reducing design risk and time to market, leading them to turn to traditional board vendors who have expertise in packaging and thermal management. PCOTS fully integrated "standard" systems enable development cost and improvements to be distributed over multiple users and programs, which drives increases in performance and economies of scale. There is also a place for both Line Replaceable Units (LRUs) as individual modules and packaged subsystems based on LRUs.

The packaged COTS advantage

In addition to ensured interoperability of the boards and mezzanine cards that populate their system, the PCOTS approach significantly simplifies system development by reducing vendor interaction to a "one-stop shop" – easing the pain that can occur when multiple third-party vendors are depended upon. The reality is that almost all deployed systems will require some unique modification. While vendors are able to offer a comprehensive range of proven, qualified systems that match most of a customer's requirements, Non Recurring Engineering (NRE) will be needed to get to the completed system. But the advantage for the customer is fast turnaround on a deployable system from an integrator who is using a common set of familiar building blocks and proven IP. PCOTS also frees the customer from the time and cost of dedicating a hardware integration team, enabling them instead to focus on the software, where true uniqueness and their value-add typically resides.

This is not to say that PCOTS is a panacea. When making this decision, customers considering PCOTS need to calculate the size of the NRE against their true overall development costs including salaries, operational costs, and so on. For example, test house costs included in a Curtiss-Wright Controls Embedded Computing (CWCEC) PCOTS project typically provide an engineer onsite for the duration of testing, and include production of the test procedures, test equipment, and test reports.

Benefits of the LRU approach

There are clear cases where it is preferable for the customer to stay with the traditional LRU acquisition approach, spec'ing and acquiring the range of needed boards and mezzanine cards from one or more vendors and handling the integration themselves. These cases include situations in which the customer's application demands a high level of programmatic, for example, requiring a dedicated team who must produce earned value reporting every month during a multiple-year time period.

Figure 1 | Avionics and vetronics PCOTS rugged systems developed from COTS components from Curtiss-Wright Controls Embedded Computing

Another case in which the advantages of LRUs outweigh PCOTS is when the application software demands a high level of custom hardware development. If the needed system cannot be built efficiently from a Size, Weight, Power, and Cost (SWaP-C) perspective using COTS components and packaged in a "COTS wrapper," then the PCOTS approach will typically not be a good fit.

The PCOTS trend on the rise

At CWCEC we are seeing increased interest from major defense contractor customers who need a PCOTS system to help them meet the challenge of increasing their card counts in space- and weight-limited ground, air, and sea vehicle environments – rather than reinvent the wheel or expend their engineering resources on packaging problems.

To get a standard PCOTS offering to meet a customer's unique application can require such modifications as custom front-panel pinouts (made simpler through a high level of modularity), and custom backplanes. One notable area of growth is the use of Air Flow Through (AFT) cooling. As interest in 6U AFT, defined by VITA 48.5, is increasing in the embedded defense and aerospace market, we are seeing a rise in the number of customers considering it. Figure 1 shows examples of CWCEC's avionics and vetronics PCOTS rugged systems developed from COTS components and tailored to specific customer requirements.

Moreover, as defense and aerospace customers discover the advantages of PCOTS – cutting schedule and performance risk while reducing their costs – the benefits of COTS will increasingly be associated not just with off-the-shelf modules but with fully integrated systems, too.

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

Do you need rugged and reliable processor power?
"Our Embedded PC Modules are for you!"

LiPPERT - for more then 25 years known for highly efficient and ruggedized products, supplying market segments as diverse as • Aerospace, • Automation, • Automotive, • Medical and • Military.

Cool LiteRunner-ECO
PC/104
Atom™ Z5xx, 1.6 up to 1.6 GHz
2 GB RAM
Mini PCI

Cool LiteRunner-86DX
PC/104
Vortex-86DX, 0.6 up to 1.0 GHz
512 B RAM
FPGA, GPS

Cool LiteRunner-LX800
PC/104,
Geode™ LX800, 500 MHz
256 MB RAM
CF socket

Cool XpressRunner-GS45
PCI/104 Express,
Core™ 2 Duo, 1.2 up to 2.26 GHz
1 GB RAM
8 MB Flash

Cool RoadRunner-945GSE
PC/104-Plus
Atom™ N270, 1.6 GHz
2 GB RAM
2GB SSD

Cool SpaceRunner-LX800
PC/104-Plus
Geode™ LX800, 500 MHz
256 MB RAM
2GB SSD

Condition Monitoring by LEMT

If your Embedded PC system must be secure, stable and running at optimum performance, constant monitoring, "Condition Monitoring", is crucial.

LEMT, being upward compatible with eAPI, provides this Condition Monitoring for your COM Express, CoreExpress®, PC/104 and EPIC boards.

For example, you can

- Continuously keep the optimum balance between Power and performance in battery-powered systems.
- Optimize maintenance of often problematic stand-alone devices. LEMT - Condition Monitoring gives you a constant overview of the device´s status and thus allows high uptimes.
- Keep repair costs and downtime low. LEMT helps here providing error codes and storing relevant data for analysis and repair.

Main Features
- Device identification
- Device status indication
- Redundant Fail Safe BIOS
- Storage of board specific data
- Measuring of current consumption
- "Secure" and "User" Flash

LEMT Architecture
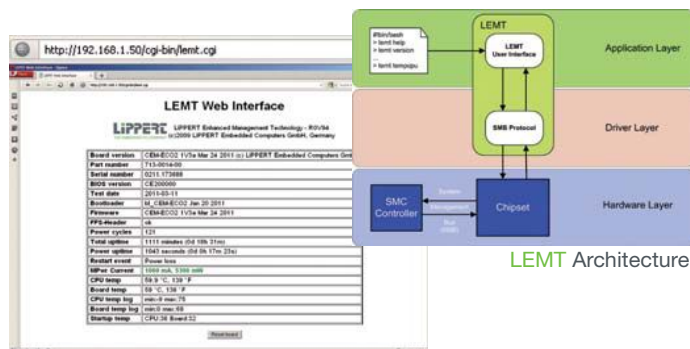
www.lippertembedded.com/lemt.html

LiPPERT Embedded Computers Inc.
2220 Northmont Parkway Suite 250
Duluth, GA 30096
Phone (770) 295 0031 · Fax (678) 417 6273

Toll Free (866) 587 8681

ussales@lippertembedded.com
www.lippertembedded.com

LiPPERT
THE EMBEDDED PC COMPANY

Daily Briefing: News Snippets

By Sharon Hess, Managing Editor

HERCULES program gains more muscle

It's not just injured soldiers in the field who need recovery and rescue; sometimes the same can be said of the combat vehicles. But field rescue for 70-ton combat vehicles such as Leopard, M1A1, and M1A2 is already at hand – and will continue to be – thanks to a recent $108 million U.S. Army Contracting Command/BAE Systems contract modification for the rescuer: the M88A2 HERCULES (Figure 1). The contract stipulates that 16 HERCULES will be supplied to the U.S. Marine Corps, while the U.S. Army will receive another 29. These latest HERCULES vehicles will be suited with U.S. government-provided remanufactured hulls, and BAE's contract work is slated for completion at the company's Aiken, SC and York, PA locales. With deliveries occurring from April 2013 through December 2013, the contract modification raises BAE Systems' HERCULES program awards to a cumulative $1.4 billion.

Figure 1 | A recent $108 million U.S. Army/BAE Systems contract covers the incarnation of 29 U.S. Army and 16 USMC M88A2 HERCULES recovery vehicles using remanufactured hulls. USMC photo by Lance Cpl. Jeremy T. Ross

U.S. Navy preps for final MCAT

Precision training is essential for military success, and a recent contract between Pathfinder Systems, Inc. and the U.S. Navy aims to achieve this. Specifically, the $6 million Small Business Innovation Research (SBIR) Phase III project stipulates that the former provides the latter with an operational prototype two of the Marine Common Aircrew Trainer (MCAT). The project builds on MCAT prototype one and consists of a baseline configuration upgrade based on Phase II simulation technologies. Work occurs at the Naval Air Station Miramar in Miramar, CA and in Arvada, CO by August 2013. The contracting activity is the Naval Air Warfare Center Training Systems Division.

Acquisitions activity: Google/Motorola head-turner, and more

It seems Acquisition Row has exploded lately – at least when it comes to the high-profile, jaw-dropping announcement that Internet search technology purveyor **Google Inc.** will acquire **Motorola Mobility Holdings, Inc.** for approximately $12.5 billion. Motorola Mobility has been and will continue as an Android licensee, already well-known for its Android-based smartphones and tablets. According to a media statement issued by the companies, the acquisition "will enable Google to supercharge the Android ecosystem and will enhance competition in mobile computing." The industry can rest assured that Android will remain open source though, Google assures, and the acquisition is expected to be finalized later this year or early next year, pending E.U., U.S., and other regulatory approvals and ratification of Motorola Mobility stockholders. In other merger and acquisition news, **Dexter Apache Holdings, Inc.,** a financial services, distribution, and manufacturing company, announced its acquisition of "100 percent of the interests of **Crystal Group, Inc.,**" suppliers of rugged COTS systems, displays, computers, networking devices, and storage devices for the military, government, and commercial markets. The acquisition's price tag was not disclosed.

OpenVPX onboard Patriot missile

The OpenVPX/VITA 65 open standard, assuring system-level interoperability for VPX/VITA 46, has found or will find its way into several of Raytheon's upgraded Patriot Air and Missile Defense Systems (Figure 2). For example, a recent $1.7 billion Raytheon Integrated Defense Systems contract with Saudi Arabia for Patriot upgrades will be fulfilled in part by Mercury Computer Systems' OpenVPX radar subsystems. Mercury has also already delivered OpenVPX radar subsystems for other Raytheon Patriot missile systems, including those deployed to Taiwan and the United Arab Emirates, Mercury reports. The recently upgraded Patriot missile system is designed to thwart threats such as tactical ballistic missiles, aircraft, UAVs, and cruise missiles.

Figure 2 | The venerable Patriot Air and Missile Defense System has been and will be suited with OpenVPX radar subsystems, per Raytheon Integrated Defense Systems/Mercury Computer Systems contracts. DoD photo by Tech. Sgt. James D. Mossman, U.S. Air Force

U.S. Navy's Hawkeye to spot more on the battlefield

The E-2 aircraft's latest iteration, the E-2D Advanced Hawkeye, will soon become more prolific in future mil ops, per a recent $761 million contract between the U.S. Navy and prime Northrop Grumman (Figure 3). The contract specifies that Northrop Grumman delivers five of its Lot 3 E-2Ds in Low-Rate Initial Production (LRIP) form. The contract also covers long lead materials to be manufactured into five Lot 4 LRIP versions of the E-2D. Not only that, the company also garnered a $34 million contract modification from the U.S. Navy in procurement of another Lot 2 LRIP E-2D. The obviously radar-equipped E-2D represents a two-generation technology leap from the previous incarnation, the Naval Air Systems Command reports, and the versatile aircraft is suited for command and control activities, battlefield strikes, and rescue operations.

Figure 3 | A recent $761 million contract and $34 million contract modification provide that Northrop Grumman supplies the U.S. Navy with more Lot 2, 3, and 4 LRIP E-2D Advanced Hawkeye aircraft. Photo courtesy of Northrop Grumman/ U.S. Navy

F/A-18E/F Super Hornets: "Start your engines!"

Some pilots of F/A-18E/F Super Hornets will soon be able to rev up their (new) engines, courtesy of a recent U.S. Navy/ General Electric Aircraft Engines contract modification. The $71 million modification calls for 18 F414-GE-400 engines and 18 F414-GE-400 engine device kits to go with them. Contract work is anticipated for completion in July 2013, and takes place at Madisonville, KY; Hooksett, NH; Evandale, OH; Dayton, OH; Lynn, MA; Rutland, VT; Terra Haute, IN; Jacksonville, FL; Bromont, Canada; Muskegon, MI; and Asheville, NC. The contracting activity is the Naval Air Systems Comand in Patuxent River, MD.

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

Figure 4 | Lockheed Martin's Squad Mission Support System (SMSS) – capable of riding in a CH-47 (pictured) or CH-53 helicopter – recently won the U.S. Army-sponsored Project Workhorse Unmanned Ground Vehicle (UGV) contest and earned the Afghanistan deployment prize. U.S. Army photo by Pfc. Phillip Adam Turner, 1st Cav. Div. Public Affairs

Lockheed's SMSS to go to Afghanistan

Having won the U.S. Army-sponsored Project Workhorse Unmanned Ground Vehicle (UGV) contest, Lockheed Martin's Squad Mission Support System (SMSS) will get the prize: Afghanistan deployment. The 11-foot-long UGV is able to transport more than 1,000 lbs of equipment for deployed soldiers on rugged soil, to reduce troops' often-100-lbs-or-more packing burdens. SMSS has been tested by the Army three times domestically already, and it's anticipated that the Army will start Afghanistan assessments in the latter part of 2011 – after training and further evaluations have occurred. An SMSS fully loaded can be easily tucked inside a CH-53 or CH-47 helicopter (Figure 4), and the UGV's sensor suite can lock on to, then follow any individual after the individual's digital 3-D profile is recognized by SMSS. GPS waypoint trails can also be followed autonomously by SMSS.

Triad of governments to share MK-99 expense

The Naval Surface Warfare Center and Omniphase Research Laboratories, Inc. recently penned a $10 million IDIQ contract for the evaluation, repair, and production of the continuous wave illuminator noise test sets dubbed MK-666. The test sets are part of the fire control system within the MK-99 (Figure 5). In addition to the U.S. Navy's 76 percent, Spain's and Australia's governments will shell out the remaining 24 percent of the contract, as part of the Foreign Military Sales Program. Contract work is slated for fulfillment in July 2014 at the company's Newport Beach, CA location.
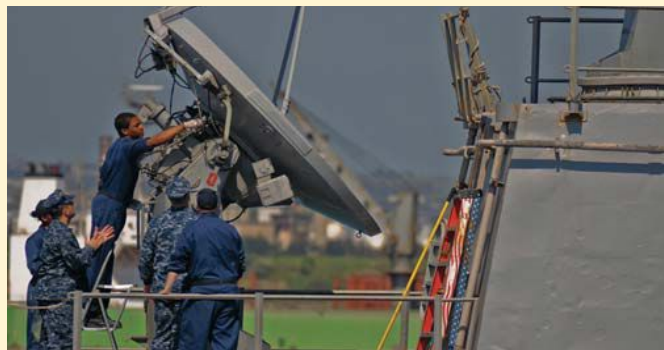
Figure 5 | The U.S. Navy and Omniphase Research Laboratories signed a $10 million contract for continuous wave illuminator noise test sets' evaluation, repair, and production for MK-99's fire control system. U.S. Navy photo by Mass Communication Specialist 2nd Class Daniel Viramontes

OpenSAF fosters high availability implementations for net-centric warfare

By Mike Houston

One of the main challenges in network-centric warfare operations is when critical systems do not achieve an effective level of "service availability," which implies a service is always available and therefore operational. A key element of ensuring system uptime is building systems with open specifications, such as the high availability middleware specifications developed by the OpenSAF community. This enables systems integrators to achieve service availability while focusing on differentiating at the application layer and through system integration.

As network-centric warfare continues to evolve, the military requires increasingly diverse and complex systems and applications. Additionally, existing systems and applications are becoming more critical because of the requirement for increased connectivity and integration to other systems. This is creating new demands for uninterrupted operation of network infrastructures. And a variety of modern military applications require this high level of service availability.

Service availability implies that a service is always available and therefore operational, regardless of hardware, software, or user fault and is often overlooked until an outage happens (Sidebar 1). The term is often used to describe scenarios where systems can remain in the field, even if they have experienced hardware or software faults. This type of solution is complex, but it also minimizes the impact of faults and automates system recovery, thereby simplifying the maintenance demand on a warfighter or operator.

Service availability is becoming increasingly important in military and aerospace because downtime of these applications can have severe consequences. For example, when a command post is engaged with warfighters in the field, continuous contact with command and control is critical for success. Losing contact could result in a loss of key information to ensure safety. Additionally, in the military embedded computing arena, one of the recent trends is placing sensors, such as video surveillance devices, on all types of manned and unmanned platforms. (Predator drones are one example of this.) An interruption or downtime in this case could result in the loss of views of activity on the

battlefield. OpenSAF (www.opensaf.org) middleware helps to ensure that high availability is achieved in such systems, to avoid these potential consequences (Sidebar 2).

Pitfalls of downtime

The costs of network downtime and the frequency of outages have been reported to be significant in a wide range of industries. Though not specific to the military and aerospace industry, a recent survey from software company Neverfail notes that 23 percent of survey respondents have had an outage for more than one full business day, while only 5 percent reported never experiencing an outage (http://extranet.neverfailgroup.com/download/Neverfail%20Survey%20Final%201.pdf). In addition, 54 percent of respondents did not know the cost of downtime. However, of those who did know the cost of downtime, 16 percent rated the hourly cost to be more than $10,000, 7 percent rated the hourly cost between $5,001 and $10,000, and the remainder rated it as less than $5,000.

Real-world examples within the military and aerospace industry related to downtime are the FAA's National Airspace Data Interchange Network (NADIN) system outages. Starting in 2008, the FAA's NADIN system, which tracks more than

---

Availability terminology

High Availability: refers to 99.999 percent (five nines) or greater of availability and typically involves recovery from faults in a system.

Service Availability: A broader concept that includes high availability, but in addition to addressing faults in a system, service availability also addresses other areas such as maintaining a system or introducing new services.

Mission critical: In the aerospace and defense industries, this term often describes systems that need to be highly available.

Sidebar 1 | Availability terminology described

1.5 million flight schedules, had experienced numerous costly service outages. These outages resulted in delays of up to six hours at as many as 100 nationwide airports at a time, largely impacting many customers across the country. The most upsetting aspect of the outages is that the FAA was unable to detect the cause of the problems in the system.

The cost of faults associated with the growing use of unmanned aerial systems in modern warfare is another increasing concern. The current concept of operations for these systems is to automatically return the vehicle to its base when an outage occurs – incurring significant operational costs and in some cases, undermining the mission and risking lives. Implementing platforms that leverage OpenSAF high availability middleware can help alleviate these issues of high availability by ensuring that all system faults are detected and isolated – and then automatically recovering the system and returning it to service.

New technologies for achieving high availability

As they seek to address the challenges like the aforementioned, Original Equipment Manufacturers (OEMs) within the military and aerospace sector face multiple challenges when providing highly available systems for customers. Intense warfighter demand, budget constraints, and industry competition in the market (Figure 1) place pressure on the system integrators and OEMs. Proprietary platforms that address high availability are costly to develop, both in terms of hours invested and expenses such as R&D. Proprietary platforms are also often complex and difficult, as well as costly to maintain and to upgrade.

Although achieving high availability can be costly and challenging, appealing COTS options that leverage open specifications are surfacing in the marketplace. Utilizing open specifications enables collaboration, extension, and a common environment for applications, making it easier and less costly to combine applications into a single platform as the system evolves. One COTS path, OpenSAF, is based on SA Forum specifications, which address the requirements of availability, reliability, and dependability for a broad range of applications. The latest OpenSAF Release (Release 4.1) is the most comprehensive implementation of SA Forum
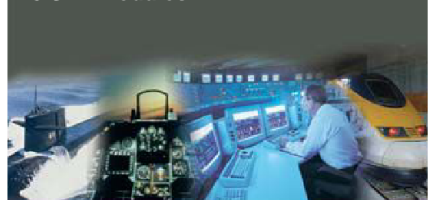
Figure 1 | OEMs face intense competition while developing and delivering highly available technologies to the military and aerospace market.

specifications available in the market, covering high availability middleware.

SA Forum specifications offer several key benefits. For example, use of open specifications helps enable system designers to leverage the commercial ecosystem for the commodity components of the system such as system management and high availability. This allows designers to focus on the differentiating application layer rather than the entire system. OpenSAF offers a modular approach to integrating high availability middleware, which allows scaling across a wide range of applications, eliminating the high costs associated with implementing a complete set of services into a platform. With a common standards-based infrastructure in place, applications can be ported, combined, and updated in a more seamless manner. This common system infrastructure streamlines overall system management, saving time and money.

SA Forum specifications define services via use of Application Programming Interfaces (APIs) for continuous hardware and software operation. At a high level, the SA Forum standards have two major programming interfaces – the Application Interface Specification (AIS) and Hardware Platform Interface (HPI).

HPI is a specification for managing a highly available hardware platform that is hardware independent and represents the platform-specific characteristics of the physical hardware in an abstract software-based model; HPI provides functions for monitoring and controlling the modeled resources.
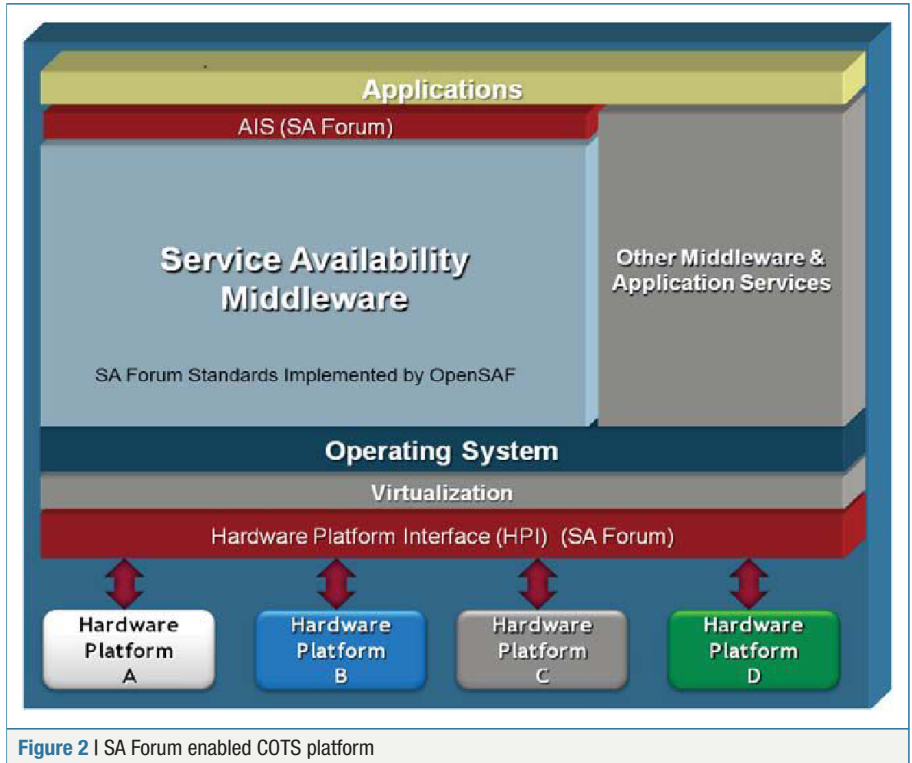
Figure 2 I SA Forum enabled COTS platform

Services from the *AIS* specification that can be implemented to ensure high availability include platform management, application management, health monitoring, fault management, cluster membership, events and messages, and checkpoint. These services manage the full range of system resources, including applications, clusters, nodes, and the platform infrastructure (hardware, operating systems, and so on).

Taken together, these AIS services and HPI provide virtually all capabilities needed for constructing a complete high availability system (Figure 2).

Service availability is mission critical in net-centricity

As adoption of network-centric defense operations continues, service availability has become increasingly mission critical in deploying military applications. It is anticipated that the requirements for five nines availability and higher will become more prevalent, along with increased usage of open standards-based commercial systems. OEMs and system integrators can build platforms with technologies such as OpenSAF high availability COTS middleware to meet these demands while also reducing costs and time-to-market. **MES**

Mike Houston is Marketing Work Group Chair of the OpenSAF Foundation and Marketing Director at GoAhead Software. OpenSAF is an open source community with projects focused on high availability middleware that align with SA Forum specifications. For more information, send email to: admin@list.opensaf.org.

PRODUCT SPOTLIGHT

OpenSAF Foundation
503-619-4109
www.opensaf.org

Eliminating the Oz factor:
How to migrate from Core 2 Duo to 2nd-Gen Core i7

By Sharon Hess, Managing Editor

Power consumption … and *thermal management* …. and *performance* … oh my. Apologies to the late *The Wizard of Oz* author L. Frank Baum for this flagrant misuse of his well-known iambic pentameter; but in the same way that *lions* … and *tigers* … and *bears* were a daunting proposition for *Oz* main characters, the challenge of migrating military systems and applications from Intel Core 2 Duo to Intel Gen 2 Core i7 processors can seem just as wide a forest for engineers to find their way in. However, a closer look at the three aforementioned key issues – in addition to pivotal software issues and the instruction sets – provides insight into easing Core 2 Duo to Intel Gen 2 Core i7 migration.

Processors' power issues: A hot topic

Considering the differences in features (and therefore instruction sets) for even just a couple of Core 2 Duo versus Intel Gen 2 Core i7 processors (Table 1), certainly design adjustments surrounding power consumption and the resultant heat are needed.

Gen-2 i7's got the power

"The [2nd generation Intel Core i7] product family draws and dissipates more power than the Core 2 Duo family. … However, it also clearly delivers highly increased performance, both in number crunching (particularly floating point) and I/O performance," says Shaun McQuaid, Senior Product Manager at Mercury Computer Systems. "These advantages greatly outweigh the relatively modest increase in power dissipation."

Accordingly, for UAVs, rugged naval systems, tactical aircraft, or other conduction- or air-cooled mil/aero applications, a 2nd Gen Core i7-2715QE delivers 134 GFLOPS per 45 W, while the Core 2 Duo might only render 18 GFLOPS per 25 W, details Alan Baldus, Intel SBC Product Marketing Manager at Curtiss-Wright. One avenue to mitigate Core i7 Gen 2's increased power is to use static control tools found in the BIOS, permitting fine-tuning of the CPU clock frequency and the quantity of active cores.

Let's just cool down here

Because increased power means more heat, a viable approach to cool things off is to use thermal simulation before migrating Core 2 Duo applications to an Intel Gen-2 i7 processor. "It may require more air than is typical [to] be flowed across the heat sink; or in the conduction case, you've [got] quite a bit of heat that must be moved to the rails. Taking a 45 W CPU in a conduction-cooled environment up to 85 °C, for example, may require a more exotic thermal solution," details David Pepper, Product Manager/Technologist for the SBC Team at GE Intelligent Platforms. (In this case, "more exotic" refers to non-aluminum materials.)

Another strategy to mitigate the heat waves of Core 2 Duo to 2nd-Gen Core i7 migration is to provide interfacing with Intel on-die digital temperature sensors, facilitating temperature control and monitoring. This is a feature that might be vendor-provided in the BSP (incarnated in the form of an API), for example, he adds.

A small sampling of available Intel Core 2 Duo versus 2nd-Gen Intel Core i7 processors

	Cores/ Threads	Clock speed	Instruction set	Lithography	Maximum TDP	Bus	Intel Turbo Boost?	Intel Hyper-Threading?	Virtualization technology?
Intel Core 2 Duo T9400	2/2	2.53 GHz	64-bit	45 nm	35 W	Bus/core ratio of 9.5	No	No	VT-x
Intel Core 2 Duo E8400	2/2	3 GHz	64-bit	45 nm	65 W	Bus/core ratio of 9	No	No	VT-x, VT-d
Intel Core i7-2715QE	4/8	2.1 GHz	64-bit	32 nm	45 W	N/A	2.0	Yes	VT-x, VT-d
Intel Core i7-2655LE	2/4	2.2 GHz	64-bit	32 nm	25 W	N/A	2.0	Yes	VT-x, VT-d
Intel Core i7-2920XM	4/8	2.5 GHz	64-bit	32 nm	55 W	N/A	2.0	Yes	VT-x, VT-d

Table 1 I Differences in Core 2 Duo versus Intel Gen 2 Core i7 processors

Performance is key

Herculean performance gains are not always a given, when migrating an application from Core 2 Duo to 2nd-Gen i7. And the real secret to increased performance is the operating system, explains Baldus. OSs enable *one* board to now house multiple applications (instead of *two* Core 2 Duo boards, in some cases), upping the performance per board.

Turbo Boost is another performance facilitator in 2nd-Gen Core i7, enabling frequency gains in a given individual core. For example, if headroom permits, a 2.2 GHz base frequency core can rise to deliver 2.9 GHz (for the i7-2655LE). But Turbo Boost can be challenging to implement. A primary consideration is sensitive timing loops, which could lack predictability unless the application was built to execute over a span of processor clock speeds.

"A timing loop that is counting on the CPU running at 2.1 GHz, for example, might be less predictable when Turbo Boost raises the frequency to 2.9 GHz. In these cases, application designers should leverage a predictable and stable clock that has some independence. Also, in the most deterministic applications, [an RTOS] such as VxWorks might be required to avoid overhead," explains Pepper.

The software calls

Though some or many of the hardware-related migration concerns inevitably involve software too, the primary, full-on software issues of the Core 2 Duo to Gen 2 Core i7 migration equation include the operating system and hardware abstraction used by the legacy application, according to Chip Downing, Senior Director of Aerospace and Defense at Wind River.

First, a good question would be: Is the application to be migrated a legacy-hardware-hosted application, and if so, does that happen via an application API or proprietary scheduler? If the answer is "yes," a labor-intensive feat is at hand. "This work will involve porting the scheduler and drivers to the new platform, modifying that software to support new capabilities like AVX [Advanced Vector Extensions] … and then performing complete testing for both the new operating environment and applications on this new proprietary platform, along with supporting a new application runtime
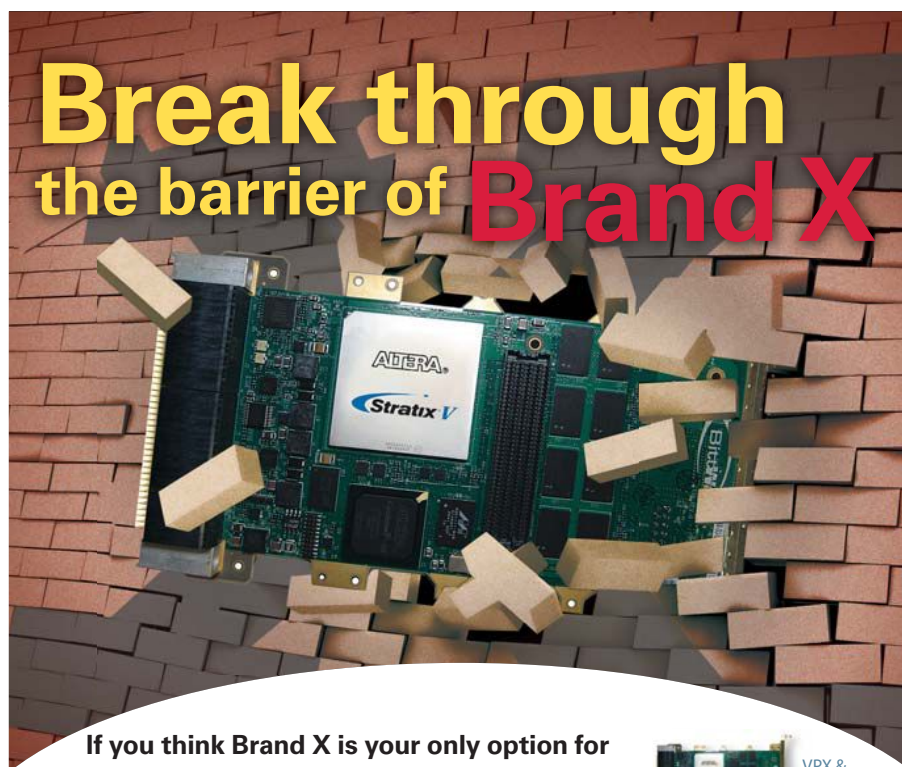
library that specifically supports AVX," Downing explains.

Luckily, these proprietary schedulers are usually not directly hardware hosted because most modern defense systems don't take full advantage of memory or CPU cores available. This paradigm ushers in a tendency toward easing migration by utilizing open standards-based software-abstraction layers. "One of the most effective methods of abstraction is the use of a commercial RTOS that abstracts away both the hardware and the I/O infrastructure, enabling the rapid

migration to new hardware," Downing says. However, even if an RTOS is used to simplify migration, if software libraries or the compiler change, a recompile and relink are necessary.

If the application does not utilize the new Core i7 features, the binaries can be moved easily to the new platform. And applications suited with POSIX APIs will benefit from the APIs' ability to facilitate immediate migration.

"The power of abstraction comes into play when using the advanced capabilities

of the Core i7 processors, like utilization of full power of multiple cores, 64-bit addressing, and virtualization. Although one could try to enable these capabilities at the application level, it is simply more efficient to let the [OS] do this heavy lifting and focus the application work on rapidly enabling requirements and competitive features," Downing asserts.

And to maximize abstraction, portability, and upgradeability and to ease multicore configuration woes on the other side, he specifically recommends using an embedded hypervisor within a virtualization scheme.

Last but not least: Instruction sets

When it comes to the Core 2 Duo versus Intel Gen 2 Core i7 processors, probably the first thing that comes to mind as far as instruction sets is AVX, a 256-bit SIMD engine providing vastly increased floating-point performance over the previous industry standard of 128 bits. Though

applications need to be scaled to harness AVX's full power, Pepper reports a "1D-FFT performance gain [of] 1.8x … when going from SSE to AVX." Scaling might include an application recompile with an AVX-supportive compiler, or perhaps an AVX function update "if SSE functions are hard coded in the application," he adds.

Onboard graphics capabilities of 2nd-gen Core i7 add 3D capabilities, with an integrated rather than a discrete controller, which changes the application's paradigm once migrated. And upgraded video, USB, and SATA are also provided in the latest Core i7 and should be considered in a migration scenario.

Meanwhile, the 2nd-Gen Core i7's aforementioned Turbo Boost gets mixed reviews when it comes to instruction sets.

"Non-deterministic behavior is frowned upon in the embedded world, so disabling Turbo Boost in the BIOS is a good practice," recommends Baldus.

However, significant performance gains are enabled with Turbo Boost, but the application must be optimized in advance for a range of clock speeds, as described previously by Pepper.

And finally, there's two-thread-per-core Hyper-Threading. Harnessing it in a migration scheme necessitates application-level analysis to figure out whether the application is able to support several processing threads, according to Pepper. Capable of facilitating application distribution for as many as eight cores, Hyper-Threading's main benefit is that performance per slot processing efficiency is much improved, adds Baldus.

Already found Oz?

Much like the characters in *The Wizard of Oz,* many system designers might already have the knowledge they need to solve their dilemma – in this case to migrate Core 2 Duo applications and systems to 2nd-Gen Core i7 platforms. But putting mettle to the metal (and binary code) is the only way to find out.   **MES**

INNOVATION IS THE FIRST UNMANNED AIRCRAFT THAT LANDS AT SEA.

Northrop Grumman X-47B
First tailless unmanned aircraft designed for autonomous carrier-based capability.

Initial Flight
Edwards AFB: 29 minutes
February 4, 2011

How does the latest breakthrough in unmanned aircraft systems go from concept to carrier deck? With help from Wind River. Our VxWorks real-time operating system was chosen by the innovators at Northrop Grumman Corporation as the software platform for their Unmanned Combat Air System-Demonstration (UCAS-D) program and by GE Aviation as the foundation for the Common Core System, the backbone of the UCAS-D computers, networks, and interfacing electronics. Building upon VxWorks' proven reliability and unmatched performance, project engineers were able to rapidly create, deploy, and maintain safety-critical control systems—all while reducing development costs and maintaining schedule integrity. Proof that when innovators work together, the sky is hardly the limit.

Please visit www.windriver.com/customers to see how Wind River customers have reached new heights.

WIND RIVER

INNOVATORS START HERE.

NASA turns to the open source community for advice

Interview with Ray O'Brien, Chief Technology Officer for Information Technology at NASA Ames Research Center

NASA's drive toward open source software (yes, you read that right) is showing up in the NASA-cofounded OpenStack open source development community and in NASA's open source software summit outreaches to the open source community. Managing Editor Sharon Hess recently caught up with Ray O'Brien, CTO for IT at NASA Ames Research Center, who talks about what's behind NASA's drive toward open source – and whether there are any resultant security concerns. Edited excerpts follow.

We received a press release saying that NASA held an open-to-the-public software summit earlier this year. Which types of open source software is NASA using – and for which programs?

O'BRIEN: NASA uses open source software everywhere and in about every language – Python, Fortran, Perl, Java, etc. And we not only use open source software, but we [design and] release [as open source] about 20 or 30 pieces of code per year – used for speeding mission design and helping us accomplish NASA science mission design, scientific missions, and engineering work. One example is GMAT [General Mission Analysis Tool], developed at the Goddard Space Flight Center to help NASA mission planners plot trajectory for space missions. It's entirely open source and written in C++; it's been used on LCROSS [Lunar CRater Observation and Sensing Satellite] and LRO [Lunar Reconnaissance Orbiter] missions, among others.

Any other NASA open source software?

O'BRIEN: Another project is Nebula, NASA's open source cloud computing initiative. Nebula contributed code helped form a new open source development community called *OpenStack* – [www.OpenStack.org], which has grown rapidly and just turned a year old; it's a really good example of a public/private partnership. Other than a standard contributor's agreement, there is no formal development commitment on the part of OpenStack community members. We just all came together to fill a gap. Anyone in the world can run [the code], or they can enhance it and NASA may also run [the code].

There's also the open source NASA World Wind application, based on Java. It's a geospatial application that provides a way for users to "tour" various portions of the Earth [http://worldwind.arc.nasa.gov/java].

Why was it based in Java?

O'BRIEN: That's just the language they selected. It goes back to why do people use open source, really. I think [developers for] every project look at the best option to address a particular requirement. Certainly, with open source, there are savings on license costs, or sometimes it's the open source

feature set, road map for future feature sets, or emphasis of the development community involved.

Something I hear frequently is, "I can actually look at the source code and I can see how something is working. And if I need to, I can actually change something to suit my needs." And sometimes if you need a feature, a vendor may commit to providing it in the future, but possibly not within a timeframe that works for you. So open source provides another option that may be the best for meeting certain requirements.

Is this reaching out to the private sector by NASA something new?

O'BRIEN: Actually NASA frequently partners with industry and other public and private entities. The twist on this is that open source *community development* is something in which NASA has not participated in the past. And Nebula [and the OpenStack community] have helped to break new ground there for NASA. The NASA Open Government team is working at the agency level with the appropriate stakeholders to develop the policy that will allow agency participation in open source community development to occur on a broader scale.

Does this OpenStack open source development community work like a standards development organization then?

O'BRIEN: Yeah, but it's more than that. They don't just talk about standards – they talk about standards, but then they create code that actually might end up becoming a standard – through adoption, not just because someone said it is a standard. Within the community, developers work on different projects, but every 4 to 6 months, they will get together at a "design summit" to plan the next release. The next design summit is October 5-7 in Boston. OpenStack members include major companies like Intel, HP, Dell, and Citrix.

Can any company join?

O'BRIEN: Yeah, any company can join, any individual can join; you can be a citizen of any country and join. So that's the beauty of the community development model – that no one cares where good ideas come from.

If absolutely anyone can join OpenStack, do you worry about security, since the software is to be used by NASA, after all?

O'BRIEN: Yeah. We worry about security all the time. But one of the nice things about open source is you can actually see how the code is constructed. You can scrutinize it to make sure it meets your security requirements. That's frequently not possible with proprietary software where you don't get the source code. And the other advantage is with the community looking at the source code, you have a much greater chance of spotting flaws in security. So, it kind of turns the argument around.

That's true, but since OpenStack and therefore Nebula are based on commonly known, open source software, would OpenStack code be more vulnerable to hacking?

O'BRIEN: Open source community development produces logic for many capabilities, but entities like NASA actually implement the open source solution themselves. Because we're in charge of implementation details and have access to the source code, we can understand how a security mechanism is coded and we can make sure we implement it correctly, so that simply having access to code does not give you a way to breach security.

What has the OpenStack development community achieved so far?

O'BRIEN: The growth of the community has been phenomenal, and it has rolled out three releases in a year. [OpenStack] is being adopted now for service delivery in public and private cloud services. NASA, through the Nebula project, is very proud to be an active member of the OpenStack development community.

How many open source community development projects is NASA involved with right now?

O'BRIEN: Except Nebula [and its involvement with OpenStack as a cofounder and active community member], NASA does not participate in open source community development. Nebula has the only waiver to do that right now.

OK. So, when did NASA's drive toward more open source begin?

O'BRIEN: A lot of it coalesced around March when we held our open source summit at the Ames Research Center – which brought together 700 people both inside and outside NASA – to discuss the open source community and how those [practices] might apply to NASA.

What did NASA gain from the summit?

O'BRIEN: Many speakers outside the government told us their views of open source, and a broader NASA community heard the benefits of open source. We used many different tools to collect the input and ended up with about 66 proposed solutions for NASA consideration, plus many views and comments on those solutions.

What were the common themes at the summit?

O'BRIEN: The number one thing we learned from non-NASA entities is that the world doesn't need another open source software license. NASA has its own open source license today, and we're advised it will be a hindrance rather than a help in getting people to join open source projects involving NASA. NASA was really innovative in 2003 and developed the NASA Open Source Agreement [NOSA] license, which really allowed us to get through some government barriers and release open source products. But NOSA has become more of a barrier than an incentive for people to participate with NASA in open source community development projects. So we need to embrace the other open systems licenses out there. The NASA Open Government team has been looking at those licenses and will be working with the appropriate agency organizations to get them approved for use by NASA.

Is licensing the biggest issue in NASA going more "open source" then, or what's the main technical problem or challenge?

O'BRIEN: You can break the major issues into policy, culture, and technology.

Regarding policy and open source community development, both licensing and the need to revise the current NASA review process for software release are the big issues.

From a culture standpoint, we need to encourage people within NASA to see open source as a legitimate way to address certain requirements. Everything doesn't have to be open source, but not everything needs to be proprietary either.

From a technology standpoint, we're pretty cutting-edge at NASA in a lot of software development. So it's just a matter of giving people permission to pursue open source community development. There are a lot of projects that want to do what Nebula did, but for whatever reason, they've never gone down the path or they've not been allowed to.

Who at NASA is doing open source software then?

O'BRIEN: All of the centers have released open source software, some more than others. Ames Research Center is probably the leader in releasing NASA open source, and Goddard releases a lot.

What's the next step in regard to open source for NASA?

O'BRIEN: One thing we're trying to do is develop one common website to list all NASA open source software at the agency level, so there's one directory, one repository. Right now Ames has its own open source Web page [http://ti.arc.nasa.gov/opensource], listing their registered open source projects.

We're also trying to update our policy internally so that open source community development is easier to pursue. And we're trying to bring the open source community within NASA together to share best practices.

Not only that, NASA is really interested in working with citizens on accomplishing its mission. So if people have ideas for open source development that NASA can work with them on, they should submit their ideas at http://open.nasa.gov. We read every submission, and we're really excited about open source because it's such a great way to engage the American public. MES

Ray O'Brien is Chief Technology Officer for Information Technology at NASA Ames Research Center, Moffett Field, Calif. Prior to this assignment, he was project manager for NASA's Nebula Cloud Computing Initiative. During his term as Project Manager, the Nebula Nova fabric controller technology was chosen as one of the two foundational components of OpenStack – the Open Source Cloud Computing Initiative. Ray can be contacted at raymond.g.obrien@nasa.gov.

GPUs shift the computing paradigm:
A 10 to 100x performance increase coming soon to a military system near you

Interview with Kevin Berce,
Business Development Manager at NVIDIA
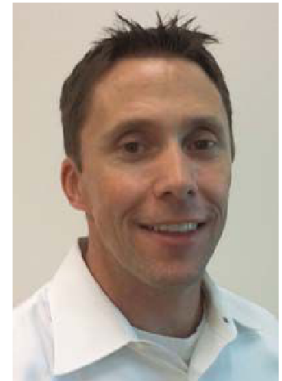
| **EDITOR'S NOTE**

GPU manufacturer NVIDIA has not only been keeping up with the times by offering processors for smartphones and tablets; it has also been enabling key shifts in the defense paradigm: Remember the days when it took four to six hours just to render one hour of UAV video? Now NVIDIA GPUs enable UAV video rendering in real-time. And NVIDIA is set to deliver an ultra-accelerated, GPU-enabled 10 to 100x performance increase for the defense industry, as a recent interview with Kevin Berce, NVIDIA Business Development Manager, reveals. Edited excerpts follow.

MIL EMBEDDED: *Let's start with a high-level overview of NVIDIA.*

BERCE: NVIDIA was founded in 1993, so we're approaching our 20th year. Our annual sales today are about $3.5 billion, and we are approaching 7,000 employees across 20 countries. The NVIDIA headquarters are in Santa Clara, Calif., and we have a global presence, especially in the Asia Pacific region, where many of our board customers are located and our processors are manufactured. We focus on engineering. We hold 1,900 patents, and continue to increase that number on a steady basis.

MIL EMBEDDED: *OK, so which types of technology is NVIDIA focusing on these days?*

BERCE: Our core heritage comes from the consumer desktop graphic space and spans gaming, high-end video editing, and other areas. Traditionally, our GPUs power everything from consumer desktops and laptops to game consoles. In recent years, NVIDIA has expanded both ends of that scale. Now we offer a range of advanced processors for mobile devices, including tablets and smartphones. Several of the new Android smartphones feature our Tegra processors, and then we also have the Tesla products where we are using the GPU to push into high-performance computing segments like government and defense. The primary need in these markets is to process and analyze large amounts of digital information extremely fast.

MIL EMBEDDED: *Which brands fall under which market segments then?*

BERCE: For mobile computing, we have Tegra; in the visual computing or the PC graphics space, then we have GeForce at the consumer end. On the professional end, which includes CAD/CAM, digital engineering, content creation for the movie industry, and much more higher-end applications, we have the Quadro brand. Finally, the Tesla product line is for the high-performance computing space.

MIL EMBEDDED: *So remind us about GPU computing.*

BERCE: GPU [Graphics Processing Unit] computing is using a CPU and a GPU together to accelerate data-intensive applications. What we have found is that the GPU has evolved from its core heritage in the graphics world to become a massively parallel processor, capable of running thousands of calculations simultaneously. This was a requirement in the graphics world, and became a technology that is very adept at dealing with large computational problems that need many processor cores to divide up the large computational workloads.

So, for example, if you have to give a CPU the problem of delivering a pizza to 20 different houses, a CPU essentially would load all the pizzas in one truck and go from one house to the next, dropping off the pizza until it completed the task. It would do this essentially in a serial fashion. If you were to give the same problem to a GPU, the GPU would take all the pizzas, divide two or three up across 20 different motorcycles, and send them all out at the same time: It would divide the task into parallel.

MIL EMBEDDED: *Let's drill down on the Tesla group for a moment. What's involved there?*

BERCE: We have the C-Class product, which is designed for workstations, and you can run multiple Tesla C-Class GPUs inside a single workstation, assuming the system has enough PCI Express slots. You can build yourself essentially a small mini supercomputer alongside your desk. The Tesla M-Class products are designed slightly differently as passively cooled and fit for server racks.

MIL EMBEDDED: *Do you have any numbers on how fast Tesla runs in supercomputers?*

BERCE: We're talking computing performance on the Petaflop level. [According to the most recent announcement of the Top 500 List (www.top500.org)], Tesla GPUs are powering three of the top five supercomputing systems in the world. Two of these systems are in China and one is in Japan.

The systems are dealing with big computational challenges. It's also worth noting that the Tsubame system is the world's Greenest Production Supercomputer, meaning that it has Petaflop-class performance; but it's also extremely power efficient, only consuming slightly over 1 megawatt. This characteristic is something uniquely enabled by GPUs.

MIL EMBEDDED: *What is the biggest power consumption? Tens of megawatts?*

BERCE: GPU-accelerated supercomputers require about half the power compared to a CPU-only supercomputer. NVIDIA strives to reduce power requirements with GPUs, thus I do not know what the largest power consumption is with other CPU-only systems.

MIL EMBEDDED: *What value are GPUs adding to the defense industry these days?*

BERCE: There are basically six verticals inside the defense space that we're focused on, where the GPU is adding value of anywhere from a 10 to 100 times performance increase: satellite imaging, video enhancement, aerodynamics/CFD, computer vision, signal processing, and electromagnetics. Our defense GPU customers include system integrators, defense contractors, and many other partners.

MIL EMBEDDED: *So what types of technical problems can NVIDIA GPUs solve, say, in UAVs or video systems?*

BERCE: One of our customers, Motion DSP, for example, has software called *Ikena ISR.* The problem they're addressing is that when video sensors get deployed in platforms like UAVs, the video acquired is of a challenging quality to obtain intelligence from.

MIL EMBEDDED: *Tell us a little about the hardware that's being used with the Ikena ISR software.*

BERCE: It's actually just a Dell workstation and inside the Dell system [running Windows], there is a CPU and one Tesla GPU, the C2070. The C2070 has 448 GPU cores, 6 GB of memory, and uses a PCI Express card that is passively or actively cooled. The value the GPU brings to the workflow is the ability to process multiple tasks in real time.

Without having a GPU in the mix, you wouldn't be able to do this in real time.

MIL EMBEDDED: *You'd have to render video offline and then get it all cleaned up.*

BERCE: Exactly, in the Motion DSP example, if the system didn't include a GPU, an hour of video would take somewhere between four and six hours to render. But with the GPU, it's rendered immediately.

MIL EMBEDDED: *What about ISR video analysis powered by a GPU?*

BERCE: IntuVision makes a software product called *Panoptes,* which is designed to enable analytics on streams of video coming in. Often in military reconnaissance or for high-end retail establishments, you have video surveillance at many different spots you want to monitor and might only have one person monitoring up to 25 video feeds. It's kind of challenging because at any given point, you can have human error and just miss something. But, having that process automated and offloaded to a computer, you know it's going to sift through all that data. With a CPU, you can roughly process 4 HD

streams at 3-5 frames per second, but we're told that to get really valuable intelligence out of it, processing up to 20 frames per second is needed. By adding a GPU into this mix, you get a 12x speedup for better object tracking and real-time notification for up to 90 HD streams.

MIL EMBEDDED: *Do most vendors buying your products – whether they're software box companies or embedded guys like GE – use a single GPU instead of multiple GPUs, typically, then?*

BERCE: The answer depends on the application they're using. They can use up to eight GPUs in a single operating system. We're seeing more and more use cases where folks are using more than one GPU. A challenge DARPA has is the desire to see how much a simulation would speed up when you add additional GPUs to the mix. That's because one of the grand challenge problems they have in computational fluid dynamics is that a single 30-second simulation takes 150,000 CPU hours. So if they can make that simulation GPU aware, they can reduce the 150,000 CPU hours substantially.

MIL EMBEDDED: *Switching gears, rocks are being thrown at NVIDIA because of the closed-naturedness of CUDA versus the freely licensed, open source OpenCL environment.*

BERCE: The CUDA programming environment was essentially what enabled NVIDIA to turn the GPU from being a regular graphics processor into a massively parallel processor that can handle the type of work that we've been talking about today. It was a far cry from the world of traditional GPGPUs. If you remember about six or seven years ago, people talked about using GPUs for computing, but they were having to translate their code from a computational language into a graphics language so that the GPU could understand it. It was very difficult. CUDA enabled developers to write in industry-standard languages to access the GPU, so that the very first version was released with CUDA C; and it was basically very similar to C. You just added a few keywords and additional instructions to change your algorithm to understand targeting a many-core processor versus fewer cores on a CPU. So there's CUDA C and now there is also CUDA C++. Third parties such as the Portland Group offer CUDA FORTAN as well. Most importantly, last week, the Portland Group introduced CUDA x86, which enables developers to compile CUDA codes to run on CPUs instead of GPUs.

So I wouldn't say that it's fair to call CUDA "closed" at this point – it supports multiple languages and open standards and can be modified to run on other architectures. CUDA will continue to be our platform for innovation; we have the ability to innovate incredibly fast. And we are able to add new features very quickly, which is something that customers in this space are very keen on us continuing. **MES**

Kevin Berce is Manager of Business Development at NVIDIA, responsible for the Defense and Intelligence business in the U.S. market. He has 15 years of experience in high-performance computing sales and business development in the U.S. Previously, Kevin was a member of SGI's Federal leadership team. He can be contacted at kberce@nvidia.com.

Rugged, prepackaged "shoeboxes" move into the mil tech fore, but what about program-specific tech?

Q&A with Paul Scanlon, General Manager of GE Intelligent Platforms Mil/Aero-Systems

EDITOR'S NOTE

In light of defense spending shifts, it's wise to keep up with the times. And GE Intelligent Platforms is doing just that: Having recently announced the expansion of its COTS Rugged Systems (CRS) rugged "shoebox" systems line to meet the mil market's latest technology shifts, GE's GM of Mil/Aero-Systems, Paul Scanlon, shares insight with Military Embedded Systems *as to where proven technology versus program-specific technology stands now – and whether it* really *matters what's inside that shoebox. Edited excerpts follow.*

MIL EMBEDDED: *GE recently announced a new line of ruggedized "shoeboxes" called the "COTS Rugged Systems" [CRS] family: off-the-shelf rugged systems designed for reduced cost and reduced time to market.*

SCANLON: Yes. About three years ago, we were primarily focused on not quite custom, but purpose-built systems, meaning a customer would bring us a specification and say, "I need this type of performance in this type of packaging to meet XYZ environments." Those were programs coming in that were funded by our customers, which are typically Tier 1 and potentially a Tier 2. However, in 2009, I had some indications from our customers, and our customers' customers, that there would be a shift in procurement and NRE funding was going to be getting scarce. So, the risk was shifting from the government to our customers in terms of deployment.

What is happening now is that we're moving from technology that is developed as part of a program, to technology that is proven and is available today – technology that can be demonstrated and that has fairly high technology readiness levels. That means reduced NRE funding, reduced time to market for our customers, and then for our customers' customer.

MIL EMBEDDED: *Even during the customization era, there was still off-the-shelf, so you're starting with a COTS set of boards and it's 80 percent COTS and 20 percent customized.*

> " ... We're moving from technology that is developed as part of a program, to technology that is proven and is available today ... "

SCANLON: It's the packaging side of things that gives the system its uniqueness, such as thermal, shock, and Size, Weight, and Power – sometimes called *SWaP*. Building a subsystem from the "ground up" is a relative term, because we start from technology that we have. We'll look at other packages that we've done, minimizing the risk around it. We are not moving away from that model – we believe that market is still there. But, we believe that there's a shift in focus to provide more proven technology. It's just the de-emphasis from a market perspective on the total customization.

Our strength is being able to pull all [the ingredients] together and implement them in a solution our customers can

count on. Our first set of CRS [COTS Rugged Systems] offerings was based on CompactPCI, PowerPC, and most recently we've offered the Intel version of CompactPCI, in 2-, 3-, and 4-slot subsystems. As we move forward, you'll see additional packaging options.

MIL EMBEDDED: *Are you using conduction-cooled 3U CompactPCI then?*

SCANLON: Yes.

MIL EMBEDDED: *Will you fit VPX or OpenVPX modules into the [CRS] model?*

SCANLON: Yes. We have a limited offering of VPX modules available today including our MAGIC1 Rugged Display Computer.

MIL EMBEDDED: *So really what we're dancing around here is VPX-based CRS products.*

SCANLON: For our CRS line, we started with CompactPCI for a couple of reasons: 1) In the marketplace today, you see a lot of need for upgrades. And 2) you see a lot of need for proven technology. And CompactPCI is a technology we've deployed quite a bit; we see growth in CompactPCI. And now VPX is just a natural follow-on to that for us.

MIL EMBEDDED: *For clarification, most of your CRS systems today are*

CompactPCI based in 2- and 3-slot pre-boxed shoeboxes. And then you're moving to VPX or OpenVPX?

SCANLON: Yes. They'll all be VPX.

MIL EMBEDDED: *What else does the future of prepackaged systems hold?*

SCANLON: We're going to higher slot counts. This year, we'll have 4- and 5-slot version subsystems and we'll also be looking at alternate cooling techniques. We'll be focusing beyond the control applications into data processing applications. And as [we] get to those data processing applications, we'll be incorporating VPX technology. We're also looking at some smaller form factor solutions. Our technology road maps [indicate] the need for smaller systems moving toward more plug-and-play applications.

MIL EMBEDDED: *Do your customers care what's inside the box?*

SCANLON: It's interesting … we've had that conversation a lot, internally, and also at the customer level, and it all really depends. If you talk at the detailed engineering level, they might have a preference as to what's inside the solution. But if you talk it up at a higher level, it's: "Does it really perform the way it needs to? Does it have the required life cycle?"

Ultimately, when we get to our customers, we end up talking about the type of product that's inside primarily from an application perspective. Our customers are taking these systems and adding value by porting their application on it, because they know which software, operating system, or target platform they've invested in, whether Intel or PowerPC. Then you start to talk about performance and there are going to be trade-offs in SWaP. Today CompactPCI can provide some of the best SWaP benefits. And as a business, we see VPX having a lot of growth in the marketplace.

MIL EMBEDDED: *How do you see CompactPCI moving forward then versus VPX?*

SCANLON: You know, if I were to draw those two on a graph, I think VPX is already there. It's gaining momentum in the marketplace, so I would say its slope

is going to be higher than CompactPCI's because of the performance of the backplane itself.

MIL EMBEDDED: *Which other form factors might be inside the GE box – and what about mixing form factors within a single box?*

SCANLON: You can look at our offerings today and see across the board the VPX solutions, the FPGA-based application or cards that come out of our sensor processing group, and the networking cards. We take all of those and integrate

them together because they're on various form factors like CompactPCI, VPX, and VME.

MIL EMBEDDED: *Do you have any [CRS] boxes available today with form factors besides CompactPCI and VPX?*

SCANLON: Not in this product offering.

MIL EMBEDDED: *What's GE's go-to-market strategy with CRS?*

SCANLON: Our strategy from a go-to-market perspective is to bring a breadth

of products to the marketplace. For example, with the announcement of the Intel version, GE Intelligent Platforms has more than 40 COTS Rugged Systems in the market.

MIL EMBEDDED: *Do you think GE Intelligent Platforms is going to sell more of these preconfigured or slightly customized CRS systems – or do you think you'll sell more boards overall (not counting the boards that go in these systems).*

SCANLON: I see a need for both. The "captive market," meaning the large tiers that do their own system/subsystem development, is a huge part of the marketplace. And they're buying the COTS boards to deploy their subsystems, so that's a huge part of our business. I see this giving them an alternative option to bring in a piece of technology that's already ready to go.

MIL EMBEDDED: *Turning to the technical side of all this, let's dive deeper into CRS. For example, how much bandwidth can the systems handle, and for which applications are they best suited?*

SCANLON: So we put CRS in two buckets here: 1) Control and mission flight computing; both have real-time processing requirements and exploit large amounts of data. And then 2) there's a different class of machine that is a subsystem used primarily in ISR applications, which is processing video and other sensor data in real time; that's what drives the bandwidth on the backplanes. Into that class of performance is where we start to gravitate toward VPX, whereas if we're getting into more I/O intensive and flight control, then that is a sweeter spot for CompactPCI.

Some of the applications we would see in that ISR space are radar, active EW, and image processing, especially video ISR where you're compressing multiple cameras and exploiting the data from those looking for targets or specific artifacts of data that lead to intelligence in real time. That's where the bandwidth is often dramatically increased.

MIL EMBEDDED: *Can you quantify the bandwidth?*

SCANLON: If we go from OpenVPX to VPX in those applications, we're going up to 10 GB on the backplane for card-to-card communication. Those are the higher-performance subsystems we're dealing with; when you start doing video over anything with multiple displays or multiple camera inputs, it gets very data-intensive quickly.

On the EW side, we're seeing it these days in the rates of 1.5 to 3 GHz, and we'll be going beyond that in the not-too-distant future. So, that certainly drives it from our perspective. On the imaging side, some image-processing platforms have a 50 megapixel focal point array.

MIL EMBEDDED: *So 50 megapixels – at which frame rate?*

SCANLON: It depends on what you're doing with the data – if you are trying to process it onboard or trying to fit it down a soda-straw downlink to the ground. It's a lot of data, so you could imagine that 30 frames per second, times the focal point array. We've seen requirements that are actually 60 or above.

MIL EMBEDDED: *750 million pixels per second is a lot and just at 15 frames; so at 30 frames, you're going to raise that to 1,500 megapixels per second.*

SCANLON: Yes, and we're starting to get more up from that, and also more distributed processing systems.

MIL EMBEDDED: *You mentioned going to 10 GbE on the backplane. Are these copper backplanes?*

SCANLON: Today, yes.

MIL EMBEDDED: *Do you see any reason to go to fiber-optic backplanes in the future?*

SCANLON: I think it's probably likely, because we're starting to hit the end stops right now on where we can go with signal integrity, both in the tracking on the boards and through the backplane connectors and also the tracking on the backplane. So we are conducting ongoing research into the feasibility of deploying optical backplanes. There are some out there in the commercial marketplace, but there's certainly nothing out there that is "rugged" just yet. We're also looking at militarized connectors that will allow us to bring optical I/O into a copper backplane.

MIL EMBEDDED: *So from the front panel of your subsystem into a backplane?*

SCANLON: Right. Within a vehicle, the high-speed Ethernet could easily be on an optical backplane, and for a variety of reasons that's very attractive. But then we have to bring that into the subsystem through a militarized connector and get it onto a backplane where we can take it back to copper within our controlled environment.

MIL EMBEDDED: *Does CRS represent a philosophical change for GE Intelligent Platforms?*

SCANLON: Let me just tell you the difference between what we're offering today in the CRS line and what we traditionally offered with our advanced vehicle computers. Advanced vehicle computers were products deployed for a specific customer and as a COTS product. They were very much configured around the original customers' requirements so they were really designed with one application in mind and designed around the space requirements.

What we've done with CRS is surveyed all the systems that we've created over the past several years and said, "Looking at the common requirements of these, here are some very stable processing and I/O products that could meet 80 or 90 percent of customer demands." And so it's a different approach in that we took all our years of market knowledge building subsystems and put together a very typical subsystem that would meet many customers' needs. That's the kind of philosophical change in our CRS subsystems. We started with CompactPCI because of market needs and we are following that with VPX subsystems to address higher performance needs. **MES**

Paul Scanlon is General Manager of Mil/Aero-Systems at GE Intelligent Platforms. He has worked at GE for more than 25 years, holding roles of increasing responsibility in engineering, sales, and marketing. Those roles have included Development Engineer, product management, and program management. In 2008, Paul took his present position of General Manager at Mil/Aero-Systems in Albuquerque, New Mexico. He can be contacted at paul.scanlon@ge.com.

Building in rad tolerance – and not as an afterthought

By Philippe Kassouf

Since no device fabrication technology is 100 percent immune to radiation effects in space, it's vital to mitigate radiation risks in space by designing in rad tolerance right from the start. "After the fact" just doesn't work.

Our mission as engineers and engineering services companies is to develop and deploy product that provides technical solutions to our customer. It is our responsibility to assess and report the risks involved with a particular technology so that mission planners understand how a product will operate in the field.

While we take great care in the design and analysis of each system to ensure the lowest possible risk, some risk is unavoidable. No matter the magnitude, it is our duty to understand, quantify, and communicate this risk. This is particularly important when it comes to mitigating radiation tolerance in products and systems designed for space. Such risk mitigation must be designed in early – at the start, even.

How to mitigate risk

Risks come in all forms: design, environmental, mission, budgetary, and so on. Each can be handled through a four-step risk-mitigation process:

1. Determine the risk
2. Analyze the risk
3. Implement the mitigation plan
4. Validate the mitigation plan

Let's look at environmental risk. If a product needs to operate at a particular temperature, we would 1) determine which aspects or portions of the system might be marginal in that region, and then 2) perform an analysis to determine the realization likelihood (expectation) of the risk by performing system- or card-level thermal analyses; 3) Mitigate the risk by designing an appropriate thermal solution such as a heat sink, fan, or heat pipe; 4) Validate the thermal solution through testing in a thermal chamber. Rinse and repeat for all other design risks.

Radiation's impact on the system

Radiation should be treated as an environment, like thermal, shock, and vibration. Independent of the source or species, radiation exists – always and everywhere – and those boards and assemblies used in such environments need to be characterized for operation. Based on the mission, the environment should be carefully characterized and the solutions tested to obtain an accurate cross-section to determine susceptibility to the environment.

And from our past experience, we have divided the solution sets for radiation capabilities into two primary segments:

1) Process and technology related
2) Design related

Process and technology related effects – or all radiation-related pathologies defined by a particular technology selection – include Lattice Direct Displacement (LDD) or Total Ionizing Dose (TID). LDD pathologies can be attributed to a neutron environment (envision a wrecking ball hitting the integrated circuit) and typically impact bipolar devices more than CMOS ones. TID, a measure of the amount of energy deposited into a material per unit mass (measured in Rad or kRad), will determine the total delivered dose to a particular chip before the onset of permanent lattice displacement damage.

Design related effects – or all radiation related pathologies that need to be mitigated by design – are not necessarily technology or process related. For example, take a flip flop (circuit with two stable states) controlling a pyrotechnic device. Independent of how the flip flop is manufactured, a catastrophic outcome can result if the device changes state "unintentionally" because of a radiation event. One remedy might be designed-in redundancy through flip-flop triplication and "majority voted" output.

Currently, no device fabrication process or technology is 100 percent immune to these types of radiation effects: The only way to mitigate them is by design.

Figure 1 | Many fields of the military and aerospace industries require reliable, fail-proof components and subsystems that ensure proper system operation.

The nature of radiation effects

The majority of design-related radiation effects fall under the umbrella of Single Event Effects (SEEs) such as:

- **Single Event Upset (SEU)** – Simply defined as a *bit flip*. This affects all bi-stable elements like flip flops, memories, microprocessors (collections of flip flops and memories), and so on. The key to an SEU is its non-permanent status, for example, the bit can be "unflipped" using Error Detection And Correction (EDAC) or scrubbing.
- **Single Event Transient (SET)** – Similar physically to an electrostatic discharge, but the charge accumulation is because of radiation.
- **Single Event Latchup (SEL)** – This occurs when a particle, heavy ion, or proton causes a parasitic structure like a thyristor or a Silicon Controlled Rectifier (SCR) within the chip to conduct from the supply to the ground. This might present an upset in operation or permanently damage the component. Note: At Aitech, we classify SEL as a "design" related effect, although it could be categorized as a process technology related effect.
- **Single Event Gate Rupture (SEGR), Single Event Snapback (SESB), and Single Event Burnout (SEB)** – Primarily manifested on CMOS, power Bipolar Junction Transistors (BJTs) and MOSFETs – these events can cause the power device to either conduct current unexpectedly, as a transient event, or permanently in a "latched" state.

Whether the event is because of a gamma ray, a proton, a neutron, an electron, or some heavy ion, to claim radiation hardness, military and aerospace as well as space system designers need to characterize the susceptibility of the design to all potential hazards of the radiation environment (Figure 1).

Of course, one can tailor the particular environment to the applicable set of tests to ensure proper operation in a given environment, hence, radiation classified as an environment. But there is no silver bullet that fits all environments optimally. Knowledge of these effects and their interactions is critical to determine how to handle them in a specific application.

Grades of radiation-tolerance process

Historically, a "rad-hard" design referred solely to total dose effects – typically 100 kRad or better – and required only proper components engineering to implement a design that could survive a given dose. Some military COTS systems might meet this level of hardness, but that's only the starting point for a total radiation hardened solution.

As feature sizes shrink, and circuits increase in performance and complexity, other radiation-related effects have started to dominate and require mitigation for a fully rad-hard solution. As a result, "rad-hard" has altered its traditional meaning, and the industry is moving towards a more stringent "radiation tolerant" or "radiation characterized" standard. Thus, the term "rad hard" is becoming increasingly misused by those who loosely extend the traditional meaning to encompass all radiation effects.

Outlined below are five grades of radiation tolerance, developed to meet the evolving radiation performance requirements and supporting the four-step risk mitigation approach. Figure 2 shows the relative complexity (and associated cost) of the different grades of radiation tolerance.

Grades of radiation tolerance

1. No radiation tolerance: Also the baseline for the engineering effort required for a given project, this provides the customer with a design that assumes no particular radiation environment. Parts and systems, qualified to military and space standards, are in fact high-reliability systems. The system might display desired radiation performance, but this is not guaranteed.

2. Total dose mitigation only: Originally known as "rad hard." The definition has evolved over time to include levels 3-5 as well. Traditionally, this step involves selecting components based on proven TID performance, and, once design is complete, optionally verifying the design by performing TID testing.

3. SEE rad-tolerant by design: As the most basic level of single event mitigation, it is purely an engineering solution to mitigate radiation in the system by designing redundancy into the system. Although Triple Modular Redundancy (TMR) is considered the "platinum standard" for achieving this goal, it does not scale very well. It is critical at this stage to balance the redundancy with system complexity to secure radiation performance, without reducing system reliability because of the increased complexity.

4. Proton SEE radiation characterized: This level of radiation tolerance is a characterization step. This level ensures the system operates in a given proton environment, and determines the system cross-section or susceptibility for a radiation event. This involves testing at radiation test facilities.

5. Heavy ion SEE radiation characterized: This is the most challenging level of characterization, because it involves the most specialized hardware and hardware configurations. Typically, this level of radiation tolerance is required for Geosynchronous Orbits (GEOs), high-inclination Low Earth Orbits (LEOs), and interplanetary missions. But it has also cropped up in some terrestrial applications.

Completion of level 4 is the point at which a product can be called "radiation characterized" or "tolerant," and Level 5 offers full radiation tolerance.

As per the risk mitigation plan mentioned earlier, until the system has been validated, there is no proof of its performance.

Hence, the grades start at Level 1 and proceed through each level to reach Level 5. The burden of each level to proceed to the next level is the engineering and testing required in implementing and verifying the performance of the product required by each level.

Relative Project Costs vs. Radiation Tolerance Level

Engineering w/o Radiation Tolerance 2%
Parts Selection for TID 8%
Heavy Ion Testing 33%
TID Testing 9%
SEE Mitigation Engineering 22%
Proton Testing 26%

- No Radiation Tolerance
- Total Dose Only. Traditional Rad-hard
- SEE Radiation Tolerant By Design
- Proton SEE Radiation Characterized
- Heavy Ion SEE Radiation Characterized
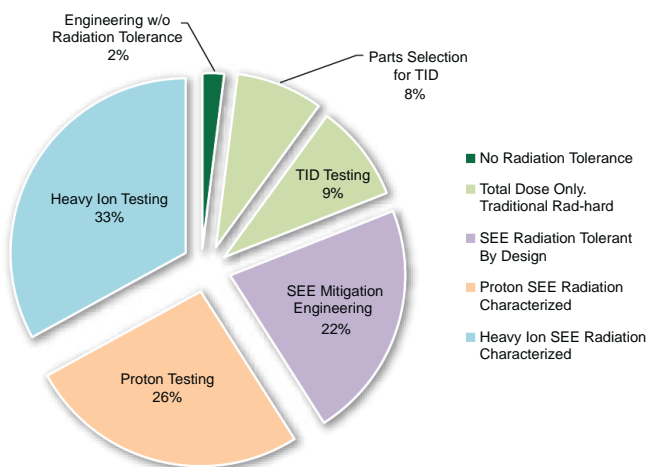
Figure 2 | Levels of rad-tolerant design requirements and their project costs

Key things to note about radiation tolerance, and its expanding attributes, are:

- Radiation data is unique. Data from one manufacturer's device does not necessarily apply to the same part type from another manufacturer.
- Radiation data is dynamic, not static. As military and aerospace as well as space systems evolve to support mission demands and as processes and feature sizes on integrated circuits shrink, radiation performance of the systems built from these chips also evolves and changes.
- Military COTS products are typically at most Level 2 tolerant, with the bulk of product in the Level 1+ regime. 1+ means that typically no more than 10 percent – comprising 2 percent for Level 1 and 8 percent for parts selection for TID – has been performed and thus accounted for.

Mitigate risk ahead of time

Radiation tolerant products must be designed to mitigate dose and SEE at the very outset, as an environment. It is not something that can be "added-on" later; radiation tolerant designs must be qualified by design and verified by test. While lot/sample testing and radiation mitigation based on technology are necessary, they're only the tip of the iceberg. **MES**

Philippe Kassouf is CTO, Space Business Sector, at Aitech. He leads a team of engineers in designing space-qualified and radiation-tolerant SBC, Gigabit Ethernet network switch, and Network Interface Card (NIC) solutions. Philippe has been involved in designing buses for micro-satellites, fault-tolerant avionics for satellites and launch vehicle avionics, and the International Space Station (ISS). He can be contacted at pkassouf@rugged.com.
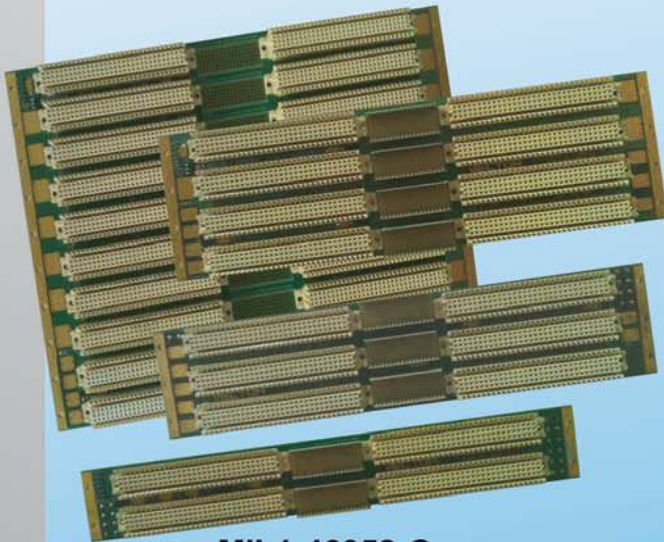
Should we put up with software that doesn't work?

By Robert Dewar

We are used to software that dismally fails. What is surprising is that we accept this as reasonable. It is time to stand up and say we are not going to put up with this anymore. There is no excuse for junk software.

To me, it is surprising how tolerant people are in general of the idea that complex software cannot be expected to work reliably. If we open the news stories for any particular week, we are almost certain to find one story about some computer error that resulted in an undesirable outcome. One such story that caught my eye recently had the headline "450 'high-risk' prisoners released in California after computer glitch."

One thing that strikes me about this headline is the use of the word "glitch." One online dictionary defines "glitch" as: "A minor malfunction, mishap, or technical problem; a snag." Oh really? Minor? What is happening here is a very typical acceptance of the idea that computer programs cannot be expected to work reliably. Anyone using a Windows-based computer is probably used to the software crashing regularly or otherwise malfunctioning, and they come to expect this level of unreliability.

This dismal level of expectations is not just something that affects non-experts. At an NYU computer science meeting, one of the faculty members teaching programming casually mentioned that all large computer programs contained serious errors. On another occasion, I listened to a presentation from an eminent

law professor from Yale Law School arguing that product liability standards needed to be modified for software because all software contains errors; so it is unreasonable to hold manufacturers to "normal" standards of product reliability. Indeed, in other areas we do expect much greater reliability. If a toy manufacturer made a mistake that resulted in dangerously high levels of arsenic in some toy, I very much doubt that the manufacturer would try to pass this off as a "glitch."

Let's take another story from recent headlines. A nefarious group of characters (perhaps associated with the Chinese government if reports are to be believed) has allegedly launched an attack on computers belonging to high government officials in the White House and elsewhere. Now in looking into the details of these attacks, they appear to principally be instances of "spear phishing," which are phishing messages specially tailored to the recipient. So the recipient might, for instance, get a message supposedly from the boss asking the recipient to look over minutes from a meeting. Recipients clicking on this attachment may find their computers compromised with dangerous malware. Now most of us react to such news by wondering how these high government officials can be so "stupid" as to click on these messages.

But is that fair? Shouldn't we perhaps be taking a closer look at the defective design and implementation of the underlying operating systems that allow such simple-minded attacks to succeed?

When safety matters

Clearly what accounts for this low level of expectation is the experience people have day to day with computers. But actually if they really knew all the interactions they have with computers, they might either be scared out of their wits, or perhaps be impressed that not all computer programs are unreliable.

In particular, whenever anyone climbs aboard a modern plane, they are trusting their lives to millions of lines of complex avionics code. Yet we have a remarkable safety record here: No life has been lost on a commercial airliner because of a bug in the implementation of avionics code. How is this achieved? By the use of rigorous techniques embodied in the DO-178B standard, we have not achieved perfect reliability (we have had some close calls), but certainly software is not the weak link in the chain when it comes to airline safety. Now I should mention in the interests of full disclosure that AdaCore is in the business of supplying tools for building software of this kind, so we have a vested interest in reliable software.

I don't apologize for that. On the contrary, I would like to see us and other companies like us succeed in convincing people to use DO-178B or similar software reliability techniques in other areas.

What about other areas? Two examples that come to mind immediately are modern automobiles and medical instruments. In the case of cars, it is on the one hand surprising that typical modern cars often have more lines of code aboard than modern commercial aircraft (Figure 1). On the other hand, it is surprising that nothing like the rigorous DO-178B standard is applied to automobile software, which is largely regarded as proprietary and not subject to the same kind of outside scrutiny. Have software bugs caused car accidents? We really don't know since manufacturers maintain a high level of secrecy. Should we worry about the future? Well, I certainly do.

Similarly for medical instrumentation, we have very complex software at work, which is also not controlled by rigorous safety standards. Have people died because of defects in such software? Yes, a number of deaths have resulted from radiation instruments delivering excessive radiation because of programming errors. It would seem that we need to tighten up the controls in this area considerably. But surprisingly, the reaction to these deaths has been somewhat muted. I think this is due to the general malaise of thinking that it is to be expected that software will have errors, so we can't get too upset about it.

Room for improvement
Going back to the avionics example, which we held up as an example showing that high reliability can be achieved, it is important to repeat that we are not perfect. At least two issues remain. First of all, we still do find bugs occasionally. So far these have not been fatal, although, as mentioned, we have had close calls – including the case of a Malaysian flight where the engines had to be turned off and restarted mid-flight because of a software defect not caught by the DO-178B certification process. No one was injured, but the plane lost 15,000 feet in altitude before the restart could be accomplished, and for sure there were a lot of scared crew and passengers. So we need to do even better.

We do have ways to further improve the process. In particular, the use of formal mathematical methods is getting much more practical. As an example, the iFacts system in England, which provides new land-based air traffic control, uses

Figure 1 | Millions of lines of code onboard: Do you trust your car?

software that has been formally proved to be free of runtime errors (the kind of thing that leads to the pervasive buffer overrun problems that plague C and C++ programs). This software is written in an Ada-based language called *SPARK,* which is particularly conducive to mathematical reasoning, and Altran-Praxis has produced a suite of tools that allows this kind of approach. Other companies are pursuing similar mathematical formal approaches. So we definitely have paths to further improvement. It's also important to note that these techniques are perfectly well applicable to COTS; there is no need to think that reliable software requires expensive customized approaches.

Secondly, we need to note that figuring out how to write highly reliable software does not guard against the situation in which someone makes an error in the specification of the software, and the resulting program faithfully does the wrong thing. An example from the avionics industry is a September 1993 accident in Poland. During the landing, the aircraft's spoilers did not deploy in time to prevent a crash in which several lives were lost. In this incident, the software was faithfully implementing a requirement that the spoilers could not be deployed unless the airplane was on the ground, as indicated by its wheels turning or having weight on the landing gear. Unfortunately, the runway was wet and the wheels were skidding and not turning, and wind shear caused a light landing. In retrospect, the specification was incorrect, and the pilot should have been able to override the wrong decision of the software.

It's impossible to guard against such errors entirely, but the process of carefully formalizing the specifications

as is needed for the certification process helps to find such errors; and for sure, eliminating normal programming bugs would go a long way to improving things. That in particular would surely have kept those 450 dangerous criminals locked up where they belong.

Don't get even, get mad
An important first step is for everyone to decide that it is not acceptable to put up with errors in programs. We need to be much more indignant when these avoidable errors occur. It's time to echo the sentiments expressed by character Howard Beale in the movie "Network" in his infamous rant and all say, "I'm mad as hell and I'm not going to take this anymore!" **MES**

Dr. Robert Dewar *is Cofounder, President, and CEO of AdaCore and has had a distinguished career as Professor of Computer Science at the Courant Institute of New York University. He has been involved with the Ada programming language since its inception and led the NYU team that developed the first validated Ada compiler. He has coauthored compilers for SPITBOL (SNOBOL), Realia COBOL for the PC (now marketed by Computer Associates), and Alsys Ada and is a principal architect of AdaCore's GNAT Ada technology. He has also written several real-time operating systems. He may be contacted at dewar@adacore.com.*

EXECUTIVE SPEAKOUT

Time and money: COTS versus traditional military system development

*By Venkataraman Prasannan,
General Manager, ATCA*

Defense agencies are using competitive bidding more and more to select project contractors. Most bidding processes where contractors must develop prototypes are winner-take-all, with only the winner compensated. During the process the agencies provide no funding, driving more aggressive schedules and increasing competition. This is difficult for defense contractors because they must invest resources and money in a project with a short development window with the goal of surviving selection rounds. As competitive bidding becomes increasingly popular, the development strategy used to win the bidding process must therefore evolve. Time and cost restraints force managers to come up with creative ways to reduce upfront development costs, decrease resources and time they spend putting together hardware, and free up resources to focus on value add of software and platform integration. With traditional military system development no longer meeting the key pain points for contractors, a COTS-based approach offers a complete solution.

Using traditional military system development, nearly every program requires custom, purpose-built electronics resulting in little sharing of development costs and intellectual property between programs. This approach has some significant drawbacks for military contractors (Figure 1). Development costs are high even before systems are tested and verified. And with minimal reuse, more redesigns are needed to correct faults, adding cost and time. One-off components are expensive due to a low rate of production, and fully custom systems take a long time to develop, ultimately reducing their useful lifetime.

Providing relief, standards-based COTS components and some simple strategies give a contractor a leg up on the rest of the pack. Compared to traditional military development,

the cost of a COTS-based system is less than half that of a fully custom system. COTS also dramatically reduces development time and cost by using standard-based platforms that are pre-integrated and pre-tested and conform to military standards.

With decreased time to market and lower investment, contractors can also meet time-critical project schedules by leveraging consulting services and fast prototyping capabilities. For example, platform specialists, such as BIOS programmers, can tune the platform, removing a large burden on contractors.

With military procurement expecting contractors to bid and develop prototypes at their own expense, they are clearly motivated to utilize these types of designs to get a head start on development. Leveraging COTS hardware can decrease resources and time even further because it can free up technical resources for software development and other value-added features that could create a competitive advantage. Also, by offering protection against obsolescence, open standards make it easier to deploy technology insertions and other performance upgrades in the field that increase the useful lifetime of systems. All of these COTS offerings help make solutions more competitive and mitigate challenges of limited funding and shortened development cycles, alleviating key pain points for contractors.

Competitive bidding is just that: highly competitive. And with reforms in the military's acquisition process, prime contractors are under more pressure to develop systems at lower costs and more quickly. Comparing traditional military system development to open standards-based systems, COTS offers unparalleled benefits, lowering investment and decreasing time to market. RadiSys is the ideal partner to succeed with a COTS-based strategy. With a deep history of technology leadership, best-in-class engineering and manufacturing expertise, and proven deployments under the most demanding conditions, RadiSys offerings dramatically reduce development time and cost, allowing contractors to meet time-critical project schedules.

Venkataraman Prasannan is responsible for the AdvancedTCA business at RadiSys Corp. He has more than 20 years of telecom and networking experience in marketing, product and business management. He has been involved in AdvancedTCA product planning and customer engagement since the inception of the industry standard.

Figure 1 | Lifetime program costs for a military system – fully custom versus COTS

RadiSys®

Network complexities be damned: The warfighter as an IT admin?

By David Kimball

Warfighters have always had a challenging job, from dealing with technologies not commonly seen outside of the DoD to frequent reassignments, with less training than they'd like. Add in the continuing shift of network technology from physical servers to today's array of physical, virtual, wireless, and cloud-based technologies across the military and it's easy to see how a military technologist can be overwhelmed. But when complexities grow, simplifying IT management for these networks can even out the odds in the warfighter technologist's favor.

There's no denying the warfighter's need for real-time information and communications on the battlefield. In today's military, a laptop can be just as important to a soldier as a weapon. Pinpointing enemy locations through UAV feeds, accessing vehicle and equipment readiness sheets, and providing on-the-ground data to a command center are just some of the many network-centric functions undertaken by soldiers during combat operations.

The downside to these expanded capabilities, however, is that as information technology evolves, so do the networks providing key services and applications to the field. Five years ago, physical servers housed the bulk of military IT resources – uncomplicated to maintain, though labor intensive. Today, these same resources are distributed among physical, virtual, and cloud-based technologies, dramatically reducing the physical footprint of the military network but driving complexity to unprecedented levels, as seen in Figure 1.

Compounding network complexity is the ever-growing training requirement of the warfighter. While command wants to shrink training requirements, advances in technology and the steady deployment of

this technology to the field require quite the opposite. Warfighting technologists now need to understand basic network troubleshooting, wireless, physical, and virtual network administration and more. Adding training for traditionally complicated technologies like network monitoring and management to a full training plate overburdens soldiers, but seems like a necessary evil given the state of military networks.

Does this mean that warfighters must act as IT administrators in addition to their roles as soldiers? In many cases, it does.

Straightforward approaches to complex requirements

The increasing complexity of IT ecosystems in the business and even civilian agency worlds is creating a shift from complex system and network monitoring technologies accessible only to trained
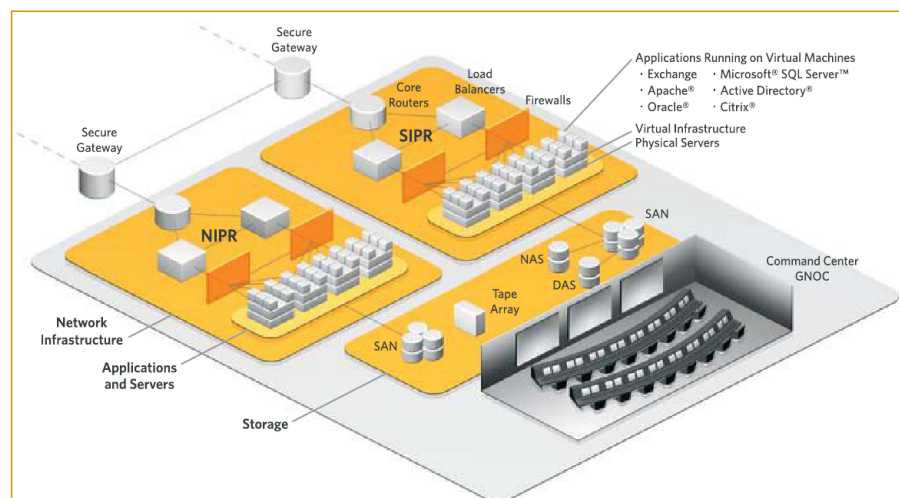
Figure 1 | Today's Armed Services network resources are distributed among physical, virtual, and cloud-based technologies, dramatically reducing the physical footprint of the military network but driving complexity to unprecedented levels.

experts, to adoption of more simplified, easy-to-learn, and intuitive technologies that allow technologists to focus on the mission and not on a network management platform.

This change in approach involves replacing traditional network management frameworks with management software that is easier to install and use. Typically, the rollout for traditional monolithic network management technologies involves:

1. A combined hardware and software installation project, which can take weeks to months based on the size of the deployment and the technology, because of integration with existing systems and hardware. (This takes even longer if the systems are geographically distributed.)

2. Customization, as most of the standard enterprise-class technologies are just "frameworks" and require lengthy (and expensive) consultation projects to meet mission needs. Essentially, these frameworks are barebones, much like a formatted hard drive, with a few simple management templates. If a unit wants management capabilities beyond the templates, which is common, customization is necessary.

3. Intensive training, on top of the training already required for warfighter technologists, to allow them to actually navigate the complex menus and esoteric interfaces of the chosen technology.

Juxtapose this with modern, streamlined network monitoring software, which is as difficult to install as your average consumer PC game. In only a few hours, monitoring software can be installed, customized, and monitor a unit's network, thanks to intuitive installation wizards and customization interfaces. Any specific training needs can be addressed via a few short YouTube videos, typically publically available online or from the vendor.

Applying easy to use, straightforward IT monitoring technologies that are efficient, intuitive, and provide immediate benefits to military networks cures many ills. First, consider the traditional network monitoring technology itself – currently, military data centers are using management platforms that are too big on features for the warzone or free/open source monitoring tools that are far too small on functionality. The problems that these approaches present are twofold:

■ First, having too many features makes the inherent complexities of managing a network almost insurmountable, especially in the field.
■ Second, a technology that is too limited does not provide enough information to actually manage a military network effectively, and often lacks the ability to scale and work seamlessly with other systems.

Finding a middle-ground technology is vital to reduce network complexity without sacrificing too many features or functionality. Ideally, the network management software will provide real-time information about network health beyond

"green light/red light" without forcing warfighter technologists to wade through dozens of unnecessary features as well as historical performance data for reporting, trending, and capacity planning. The technology should also be able to scale with the network, up or down, depending on need: If a FOB expands to encompass more than a single unit, the system must scale up, and conversely scale down if a base is decommissioned or downsized.

The network's physical size matters
Simplified network management technology goes beyond just reducing the digital complexities of military networks; it also greatly aids in reducing the physical

requirements of networks. Larger IT systems need more power and cooling, which require a greater amount of fuel deliveries and power infrastructure. Reducing energy consumption is not just a "green initiative" for the military, but allows for greater mobility and a smaller logistics tail, particularly in operating theaters where a data center could be a 2' x 2' box on the back of a Humvee. DoD initiatives are also underway to reduce physical footprints by using powerful laptops as servers and leveraging virtualization to reduce the server count; virtual servers that host 10 or more virtual machines are increasingly common and, in some cases, those numbers can be increased by a factor of 10 or more.

Deploying streamlined, modern network management software that is already "virtualization-aware" allows military data centers to harness the cloud and virtual environments without requiring warfighters to learn a new monitoring system. However, with a traditional network monitoring paradigm, not only is there a training requirement for each system implemented, but each system desegregates management data and makes issue resolution and information correlation even harder. But by using software that manages as much infrastructure as possible within one single pane of glass, soldiers can monitor virtual and cloud-based resources just as they would a physical box, without having to learn new, specific skills.

This leads to the next point of how simplifying network management technology addresses complexity and reduces the training required.

Simplify to reduce the training requirement

One of the primary problems caused by complex military networks is actually related to a larger problem experienced by the whole of the DoD: warfighter training requirements. Not only are soldiers expected to learn combat, survival, and tactical skills prior to deployment, now they must learn the complexities of network management and monitoring just to use the tools issued to them, hence the "warfighter as an IT admin."

When the management of these networks is simplified, however, overall training requirements follow in reduced complexity. Instead of learning how to navigate a complicated monitoring system with
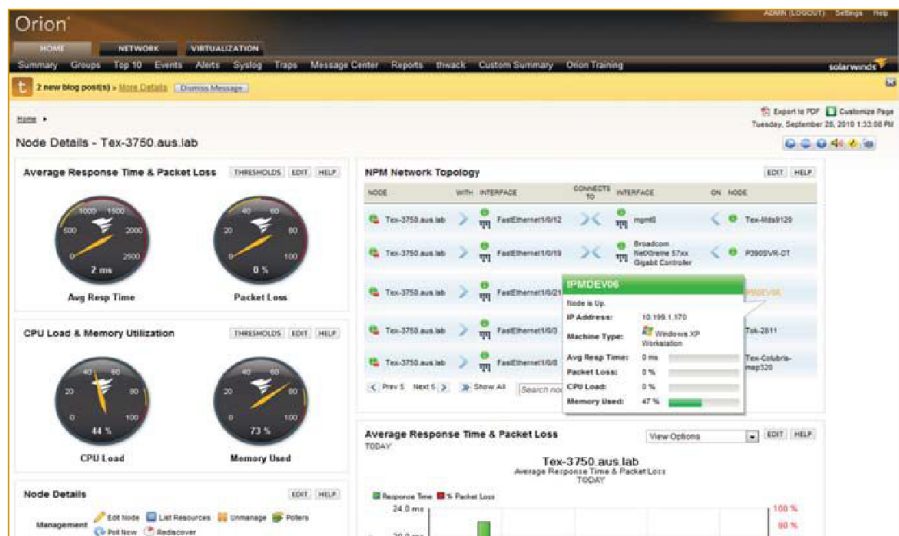
Figure 2 | The Orion Network Performance Monitor is an example of a streamlined, simplified network management software

more features than needed, soldiers can gain a few simple skills that apply across monitoring/management challenges. Much like issuing iPads to soldiers in Afghanistan in lieu of a ruggedized laptop, simplified management software applies to a much wider base of warfighters and provides the features most needed to accomplish the mission. As an added benefit, this software can run off COTS hardware and operating systems, including open source, making it flexibility enough to fit nearly every unit's operations.

Additionally, the threat of budget cuts must also be considered. It shouldn't surprise project leads when they suddenly have fewer experts to operate clumsy legacy IT monitoring systems. Without the right team mix to prop up these sputtering dinosaurs, leads must expand their use of efficient IT monitoring applications, to not only reduce the need for extensive training on legacy systems but also to cut the costs of maintenance renewals.

Simplifying for a complex tomorrow

Network complexities are not going away and will only become more severe, particularly in military networks, unless addressed head-on by network management technology. Leveraging modern IT management software, such as that offered by SolarWinds, that supports both mature and emerging technologies helps prepare military IT teams for rapid technology insertions. SolarWinds' Orion Network Performance Monitor is an example of a streamlined, simplified network management software (Figure 2).

When such IT management systems are also intuitive to use, configure, and maintain for even the most junior IT staff members, the mission can then receive the benefits of the latest technology advances while reducing complex training and staffing requirements necessitated by legacy IT management applications. Monitoring the most modern technologies requires the most modern management applications, but "modern" no longer means complicated. This is especially true for those soldiers, the "warfighter IT admins," who are responsible for managing the technology that keeps their teams running. **MES**

David Kimball is VP of Public Sector at SolarWinds. He has more than 20 years of IT sales and senior management experience, with more than 10 years of experience focused on the federal government market. Prior to SolarWinds, he was Managing Director of Federal Sales for Cenzic and the Regional Vice President of Federal/Enterprise for SupportSoft (acquired by Consona). In addition to sales leadership positions with Intellitactics and NetIQ, David also designed and built a strong federal sales organization at Remedy Corporation. He can be contacted at david.kimball@solarwinds.com.

Designing rugged, multifunctional HPA controllers to prevent system damage

By Mario Razo and Meppalli Shandas

High Power Amplifiers (HPAs) are critical in microwave systems commonly found in Electronic Countermeasures (ECM), Electronic Warfare (EW) simulators, radar, and communications links used by the military. Unfortunately, integral Vacuum Electron Device (VED) requirements introduce hostile environments for HPAs. However, a rugged, multifunctional, microcontroller FPGA-based controller can solve this challenge.

High Power Amplifiers (HPAs) are the backbone of most microwave systems used for military applications such as radar, Electronic Countermeasures (ECM), communication systems, and Electronic Warfare (EW) simulators. HPAs using Traveling Wave Tubes (TWTs) come in two categories. Both versions use Vacuum Electron Devices (VEDs), TWTs, Klystrons, and Gyrotrons to amplify the modulated RF waveforms given at the input to the desired power level before feeding to the radiating element.

The difference is that one version of the HPA is the microwave power amplifier with all RF input and output parts, power amplifying TWT (or other VEDs), digital interface and protection circuits, and the power supply all integrated into one assembly. The other version is the compact Microwave Power Module (MPM), which uses a miniature version of the TWT and a solid-state driver amplifier integrated with a densely packaged power supply (Sidebar 1). Being a very compact module, the MPM does not have all functions within the assembly and can produce RF power in the range of about 100 W CW or 1,000 W peak pulse power. The VEDs require high operating voltages of 5 to 25 kVdc and proper switching sequences and protection circuits.

Winning the battle with wide-bandwidth, high-power microwave functionality

Military systems such as radar, Electronic Countermeasures (ECM), and communication links compete in a modern Electronic Warfare (EW) game of "cat and mouse" with the enemy. To win, design engineers must create wide-bandwidth, high-speed secure systems that are invulnerable to countermeasures and electronic attacks and will perform reliably in challenging battle situations.

In the lower-frequency domain below 2 GHz, designers are developing reconfigurable systems for specific tactical needs. Highly integrated hardware, unique algorithms, and exotic waveforms ensure that transmitted data is secure and impossible for opposing forces to decipher. But in the higher microwave frequencies, wide-bandwidth, high-power amplifying devices aren't as readily available. The most commonly used device is the Traveling Wave Tube (TWT). For example, a TWT capable of operating in the 6 to 18 GHz (or even 2 to 18 GHz) range and delivering 200 W of output power can provide instantaneous bandwidth over the entire range. Operating in the C, X, and Ku band frequencies, TWTs cover all multiband applications such as ECM, radar, and data links. These are generally known as High Power Amplifiers (HPAs). A lower-power version of the same HPA function can be accomplished by a Microwave Power Module (MPM).

Higher-power radars using Active and Passive Electronically Scanned Arrays (AESA, PESA) are being modified to perform electronic attack function. These are extremely agile systems, but because of performance limitations of the Transmit/Receive (T/R) modules using solid-state technology, these systems can cover only a small bandwidth.

For platforms with limited prime power and severe restrictions on size and weight – such as Unmanned Aerial Systems (UASs) used for Intelligence, Surveillance, and Reconnaissance (ISR) functions – MPM is a good solution. Basically using a miniature version of the TWT, the MPM can meet all tactical needs for bandwidth, speed, and security. This compact, high-power amplifier can be used for multiple roles such radar, ECM, and data links. When configured as an embedded transmitter, an MPM will solve most high-power amplification problems.

Sidebar 1 | An inside look: How microwave systems solve high-power electronic attacks

To keep the HPA controller functioning is vital for continued effective systems operation. Thus, to prevent damage to costly parts and the system, ongoing monitoring and status indication of every critical parameter are essential. For example, HPA controllers must monitor HPA health by measuring critical parameters such as the cathode voltage, current, body temperature, and so on. Controllers must also provide protection to prevent damage, in addition to offering interface between the HPA and host system, executing housekeeping functions, and functioning as the Built-In Test (BIT) system for the HPA.

In addition, HPA controllers must be able to withstand a "hostile" environment, as they reside in an area of the system where they can be subjected to high voltages up to 25 kVdc, high-energy switching up to 10 Joules, and other severe operational stresses such as extreme temperatures (-55 °C to +125 °C) and high vibration levels of up to 20 G rms. This is further complicated by the high levels of switching currents in the order of 100 A per microsecond and short-circuit currents of thousands of Amps. If such currents were to be going through any of the VEDs, it will destroy these expensive devices beyond repair.

The design challenge is to provide suitable protection and integrate BIT functionality and continuous monitoring into a microcontroller FPGA-based embedded controller because of the HPA's "hostile" environment. Design engineers must consider the "hostile" environment factors prior to establishing the appropriate architecture for HPA embedded controllers. Figure 1 illustrates a basic HPA system architecture. Three major modules form an HPA system: a VED, a controller, and a high-voltage power supply that provides the required voltages to different elements of the VED such as the heater, cathode, collectors, and the electron beam. The HPA controller offers control and protection for the entire system.

Protecting HPA systems' critical components

Every HPA system requires control and protection circuitry to protect the VED and other critical system components from adverse conditions and possible failures that could lead to major system damage and considerable economic impact. There are several factors to consider when designing high-performance HPA embedded controllers. For example, the high operating voltages (5-25 kVdc) required by an HPA system can cause
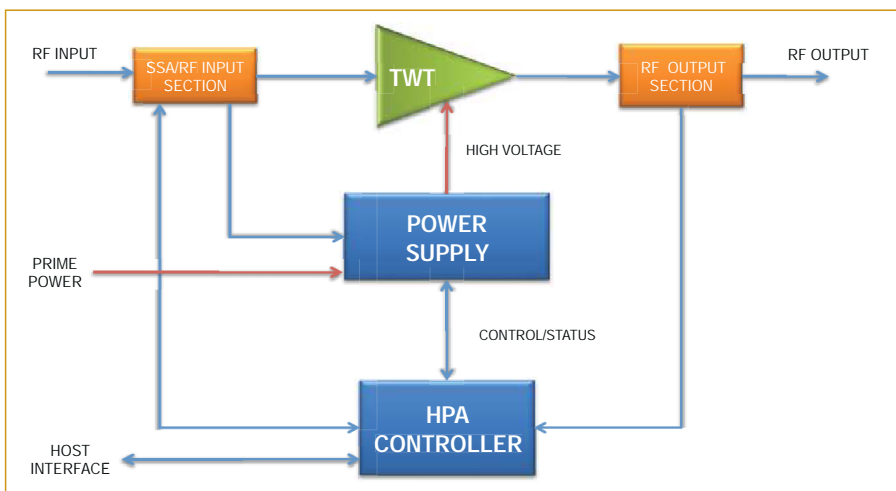
Figure 1 | The basic HPA system includes a VED, power supply, and HPA controller.

severe damage to sensitive components such as microprocessors, microcontrollers, FPGAs, memory devices, Analog-to-Digital Converters (ADCs), and other critical components. Extreme protection is required to keep these devices safe and active under any circumstance, especially when high voltage and switching current spikes occur. Current spikes of HPA systems can vary from a few Amps peak up to hundreds of Amps peak, depending on the VEDs' requirements.

The hostile environment surrounding the embedded controller can also cause false alarm failures or unexpected behavior. Switching noise produced by the power supply switchers can produce extremely sharp noise spikes of nanoseconds duration, which can occupy more than a few MHz of bandwidth. These can corrupt critical signals of the HPA system such as Serial Peripheral Interface (SPI), Inter-Integrated Circuit (I^2C), and serial data lines.

The HPAs need to meet various conditions such as time delay required by VEDs (180 seconds typical), operate command from the host interface, and Focus Electrode (FE)/grid enable and beam control (-1100 Vdc to +500 Vdc) prior to amplifying any given RF input signal of 2 to 40 GHz. The HPA controller must keep control over these signals and execute the commands in the right sequence under any condition to assure safety and proper functionality of the HPA system. Improper switching sequence of these signals could damage the VEDs, accidentally transmit RF output power (in some cases 10 kW peak), or harm personnel. Implementation of control signals with safe-state configuration is necessary to achieve a suitable and safe HPA controller. Safe-state signals will protect the VEDs from being active during initial power-up, standby, or when transmission is not desired. The integration of microcontrollers and FPGAs into HPA controllers empowers HPA systems by providing high-performance protection and extensive functionality, as will be discussed later.

"Universal" controllers satisfy next-gen systems
Modern mission-critical systems require HPAs that meet all performance requirements, provide 100 percent availability for the mission duration, and act as multifunctional HPA controllers capable of withstanding high-stress environments while effectively performing essential
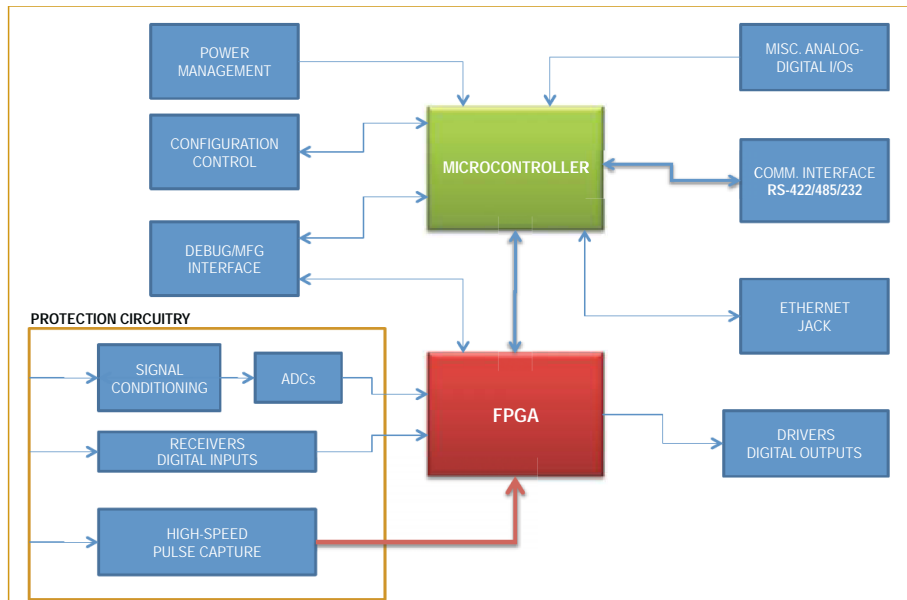
functions. A high-performance "universal" HPA embedded controller with an input power of +15 Vdc, 5 W, and dimensions of 6.0" x 2.5" x 0.90" can meet these requirements for next-generation military airborne systems. Figure 2 depicts its basic block diagram and hardware architecture.

The main function of this embedded HPA controller is to protect the critical elements of the system by monitoring the diverse parameters of VEDs such as Helix current up to 600 mA, cathode voltage from 5 kV to 25 kV, over-temperature (typical 125 ºC), and output reflective power up to 100 W average. These parameters vary from VED to VED, and the embedded controller offers flexible protection limits. It also monitors other critical parameters of the HPA such as input current, input voltage, and temperature of the high-voltage power supply. Additionally, the controller provides built-in test features through different serial communication protocols such as RS-422, RS-485, or RS-232. This function enables the host interface to continuously retrieve (automatically or upon request) the HPA system's status. The HPA controller provides status between 1 and 2 seconds after application of power.

The embedded HPA controller can receive commands such as Model Identification Query, Status Query, Operate, and Standby, for example, from the host interface via the serial interface and execute

these commands almost immediately (in less than 20 ms) after reception and acknowledgement.

Using a microcontroller, FPGA to protect the HPA
The microcontroller and FPGA components are the main core of the embedded HPA controller that protects the EW, ECM, radar, or other military system electronics from damage; these controllers perform the majority of the tasks with aid from the protection circuitry section and other miscellaneous devices, as follows:

■ The main functions of the microcontroller are overall supervision, host communication, self-test, and technician assistance features via RS-232. The microcontroller also verifies board/system configuration, communication with a remote panel, and communication with other controllers, and conducts periodic verification of FPGA configuration to guard against single-event troubles.

■ The protection circuitry continuously monitors the HPA system's critical parameters by collecting feedback from various HPA elements. Such feedback could be temperature readings (-55 °C to 150 °C), cathode voltage, helix current, reflected output power from the VED, input under-voltage, and more. The input gate signal (pulse systems only) is detected by a pulse measurement function inside the FPGA to protect the VED from over-duty of

3 to 30 percent, over-pulse width of 2 μs to 300μs, and over-frequency of 200 Hz to 100 KHz. The protection circuitry also incorporates a high-speed pulse capture that accepts RF detector signals. These signals are conditioned, scaled, and digitalized, and the FPGA transmits the digitalized values through the serial interface.

The FPGA enhances the HPA controller by providing an extensive set of features to make it more robust and multifunctional. The FPGA architecture includes configurable functions via tailored logic loads, polarity, and mask registers to allow a common logic load for multiple HPA systems.

The FPGA provides control of a 12-bit ADC and compares configurable thresholds against digitized voltage inputs such as cathode voltage (scaled down from 5-25 kV to less than 10 Vdc) and detected RF output power (-80mV) from the VEDs (typically amplified and inverted from -80 mV to +10 Vdc). Additionally, the FPGA captures and analyzes digitized data, performs gate-pulse measure-ments (2-300 μs with frequencies up to 100 KHz), and provides voltage-to-power conversion using a lookup table, in addition to peak detection.

As a safety feature, the FPGA is configured to keep the I/O signals in safe-state; the microcontroller is capable of detecting any FPGA malfunction on the next polling cycle and taking appropriate action.

HPA controllers for today's military systems

While HPAs are essential in most micro-wave systems used for military applications such as ECM, radar, communication, and Electronic Warfare simulators, it can be challenging to integrate the necessary protective functions into the HPA's embedded controller. The hostile environment introduced by HPA systems because of the VED's requirements is inevitable, and it compromises the integrity of HPA embedded controllers. Design engineers are obligated to live with this matter and must protect the embedded controller. dB Control has developed such an embedded HPA controller, which is microcontroller FPGA-based, noise immune, multi-functional, and firmware configurable, minimizing the need for design changes and harnessing FPGAs' and microcontrollers' high performance and extensive functionality to protect the system. **MES**

Mario Razo is a Design Engineer at dB Control, working with high-power amplifiers and microwave power modules for military and commercial customers. Email him at mrazo@dbcontrol.com.

Meppalli Shandas is dB Control's VP of Technology and Business Development. He has 40 years of design experience with microwave hardware for Electronic Warfare, radar, and communication systems. Email him at mshandas@dbcontrol.com.

Combat fire support
and situational awareness?
There's a smartphone app for that

By Michael Howard

U.S. Army photo by Heather Vann

Mobile applications are just starting to earn their way into the "trusted" status of the military IT toolbox. But given the growing use of smartphones on the battlefield, can a smartphone do more than just monitor a network or provide secure chat? Could it actually heighten situational awareness to help save warfighter lives or even provide aid in combat fire scenarios? The answer is a resounding "yes," and a Composable Handheld Android Platform (CHAP) comprising a Call For Fire (CFF) Mission Thread is key.

Despite technological advances across the board, the Department of Defense continues to grapple with a familiar capability gap – situational awareness, particularly the availability to deploy small, dispersed, mobile units. While situational awareness tools are now heavily automated, the route that many DoD technologies are taking, this automation emerged in the matrix of legacy command and control architecture with thick clients, voice communications, authoritative data sources, and centralized command centers. However, doctrine and operational practice have now outgrown such an architecture. Small, highly mobile, increasingly autonomous units (and even individuals) are intended to operate in a net-centric fashion, but legacy issues of network infrastructure, heavy systems, and lack of access to data continue to limit mobile situational awareness.

Combat fire support poses a particularly tough challenge for mobile units. Its stringent requirements for timeliness and accuracy and inherent lethality as well as dependence on an accurate operational picture suggest that if a system works well and reliably for fire support, it should also lend itself to other domains. Such a mobile combat fire support system now exists for military use, the Composable Handheld Android Platform (CHAP), which comprises a Call For Fire (CFF) Mission Thread, Infrastructure as a Service (IaaS), Software as a Service (SaaS), a Web portal, and an Android or iPhone smartphone application and associated services. The following discussion examines the different components that enable a CFF to be executed from the CHAP, from functionality of the device to the makeup of the systems behind CHAP.

Eyes on situational awareness

Complete situational awareness comes from eyes on a target (for example, video feeds from UASs), Blue Force and Red Force locations and descriptions, current and accurate imagery/maps, and local area intelligence. With the aforementioned technology unified into a single system, CHAP provides an observer with necessary intelligence, as well as the ability to modify, delete, or create tracks (actual moving targets) directly from the mobile device. Fire support can now be called in through a Droid or iPhone application, simply by confirming the target via the touch-screen; this smartphone-based support also gives the warfighter access to other intelligence via data services functionality, including chatting with other units currently connected to CHAP in the cloud. A screenshot of the soldier's view can be seen in Figure 1.

When a target is created, it is reported through the cloud, processed through a Complex Event Processing (CEP) engine, and sent to the secure commander portal in real time for review and authorization. Via the portal, the operation commander can view an entire Common Operational Picture (COP), which runs as SaaS within the IaaS. Additionally, any authorized

Figure 1 | Droid soldier view – Call For Fire (CFF) Android application
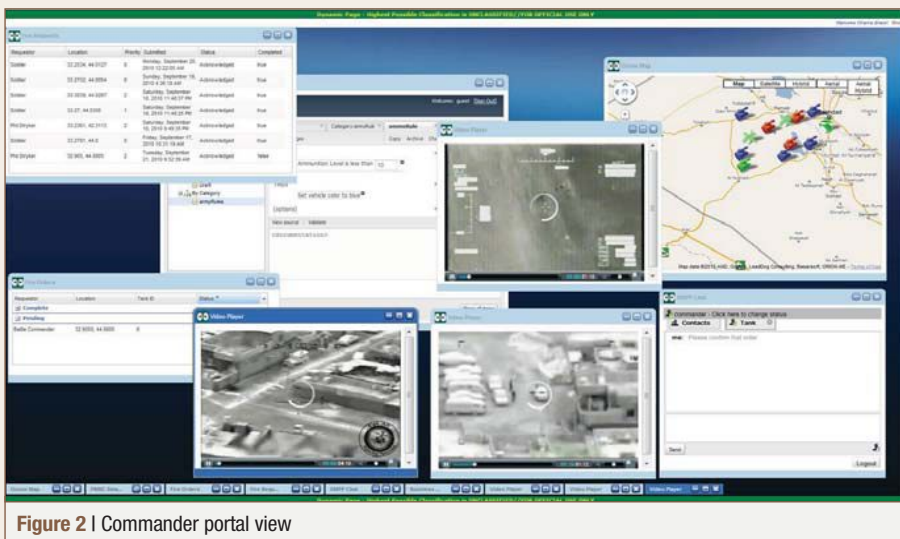
Figure 2 | Commander portal view

users in the portal can create a customized view, or User Defined Operational Picture (UDOP). The user can select from a library of widgets (such as Google Earth and NASA World Wind) to build their UDOP, which can be seen in the screenshot of the Commander's view (Figure 2).

The prime question remains, however: What does each disparate piece of technology bring to the table, ultimately enabling mobile situational awareness for the warfighter? We'll start by examining the CFF Mission Thread.

Live fire – CFF Mission Thread
The CFF Mission Thread is a series of documented processes, involving dozens of parallel systems and applications as well as high command review built around an observer (for example, a warfighter) using a smartphone to call for fire. This is done simply by the user holding his/her finger on the target for a few

milliseconds until the command is transmitted to the cloud for analysis, against a target, whether an enemy vehicle, strongpoint, building, or unit. The Mission Thread is the backbone of CHAP, allowing the warfighter to call for, and receive, fire support in a hostile situation. For the Mission Thread to function as needed, however, several criteria must be met.

First and foremost, the observer must be connected to the military network and the supporting applications, like the Address Book, must be online and functional. Next, the operation commander's criteria and Fire Support Coordinating Measures (FSCMs) must be readily established guidelines for the type of fire support to be received: how this incoming fire will affect nearby operations and whether or not air support is to be received. Finally, the observer must have a laser (used to "paint" the target/targets), and there must be weapons positioned in the area of operations to actually fire on the target.

When these criteria are met, fire support can be received. As for the type of fire support, Precision Guided Munitions (PGMs) must be employed, which provide near-pinpoint accuracy to the observed target. In this case, PGMs can include supporting fire from howitzers, tanks, unmanned aerial systems (like the Predator drone), and attack helicopters.

To prevent unauthorized use or interception, the Droid or iPhone used by the observer could be equipped with a Type I Encryption Sleeve. Currently undergoing prototype testing with NSA, this is similar to a standard smartphone case or iPhone "powerpack" that locks around the outside of the smartphone, providing Type I encryption capabilities to a typically unsecured device.

Fire from the cloud
CHAP relies heavily on the cloud for its flexibility and overall functionality to the warfighter, primarily via IaaS and SaaS. IaaS renders the virtual environment transparent to the end user, allowing updates to the applications (such as the Droid app, services, business rules engine, portal, and widgets) and the underlying infrastructure to occur without procuring physical memory, storage, or processing power. Meanwhile, SaaS enables easy deployment of COTS software for additional functionality within the technology. Enterprise services provide a foundation for interoperability among the capabilities deployed within the cloud environment and consist of eight separate services:

1. Discovery Service
2. Messaging Service supporting AMQP and Java Message Service (JMS)
3. Orchestration Service
4. Real-time Collaboration Service
5. User-facing Services
6. Security Services
7. Mediation Services
8. Enterprise Service Management (ESM)

Along with the aforementioned enterprise services, the Department of Defense's Universal CORE (UCore) data model is also used, allowing the simulated data services running in the cloud to talk to each other. Additionally, UCore in this application is extended to support Track Metadata, Video MetaData, and Variable Message Framework (VMF) messages for position location, along with Keyhole Markup Language to support Google Earth Integration.

The simulated data services include multiple intelligence systems, including sensor data from remote monitoring units, video feeds from UASs and aircraft, maps, and other imagery. Most notably, the Global Command and Control System-Joint (GCCS-J), a DoD legacy system, is also run in the cloud as a simulated data source. To do this, a tracks service (an application/messaging service to transmit tracks/target information) was written on top of GCCS-J, allowing it to push its data into the cloud for use by the other cloud-based services.

For the final step in the cloud, a Universal Description Discovery and Integration (UDDI) registry stores all the services and their associated data models for CHAP, which allows the disparate systems making up CHAP to find the appropriate service when needed for a specific event, like an acquired target or a request for fire support. A CEP engine then processes these multiple events and messages to determine which should be acted upon, in conjunction with a Business Rule

Management System (BRMS) that defines, deploys, and executes business rules for the system.

As an example, the CEP engine has a rule to easily let commanders see when a tank or other mobile firing platform is running low on ammo. More specifically, when a tank's ammo threshold drops below 20, the affected tank turns blue, allowing commanders to easily see which tanks are friendly (blue), hostile (red), or low on ammo/resources (green). The Commander Portal contains a widget to view this on the Map for SA.

Forging ahead under fire

This technology was co-developed by QinetiQ North America and Red Hat in their internal research and development program. The current technology source code is available via projects within the DoD's Forge.mil, including an Android project established during this technology's development. The system described herein was developed to demonstrate the art of the possible for the current

tactical and enterprise environment. The technology is being refined for deployment to warfighters who need near-real-time situational awareness – and that, of course, is every warfighter whose boots are on the ground. **MES**

Michael Howard is the Vice President of Advanced Enterprise System Solutions for QinetiQ North America's Systems Engineering Group. He has been recognized throughout the DoD for his efforts in Service Oriented Architectures (SOAs), Master Data Management (MDM), System of Systems (SoS) Engineering, and SoS integration. He holds a BS in Computer Engineering from the University of South Carolina. Email: Michael.Howard@qinetiq-na.com.

QinetiQ North America
843-740-7456
www.qinetiq-na.com
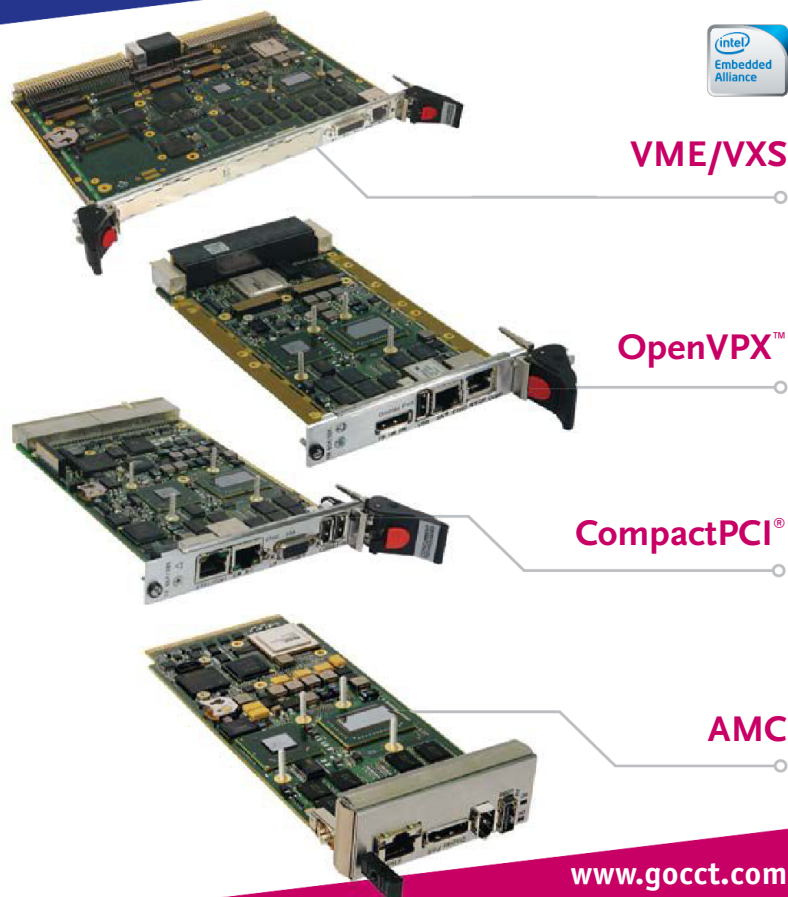
Smartphone encryption app helps sensitive information get more secure

Interview with Ian Meakin, VP of Marketing at Cellcrypt, Inc.

Off-the-shelf iPhone, Blackberry, Nokia, or Android-based smartphones are simply everywhere and used by most everyone these days. And off-the-shelf smartphones are undoubtedly used by military officials, foreign dignitaries, and federal agents (or their staff members) hustling to schedule meetings. In the process, they reveal "sensitive but unclassified information" in the clear, including where and when to meet, with whom they will meet, travel plans, and more. Though this information seems rather mundane, if it is intercepted by hackers, it can present a high risk to national or personal security. But CellCrypt, Inc. says they've developed a remedy: their secure voice calling app for smartphones (including secure text for Blackberry), used by the military and utilizing NSA Suite B and FIPS 140-2 crypto to secure sensitive but unclassified information. In the following interview, Managing Editor Sharon Hess got a behind-the-scenes look at how the application works technically, and found out whether it can protect the next level up: classified information. Edited excerpts follow.

> **Tell us a little about Cellcrypt, its technology focus, number of employees, locations, and so on.**

MEAKIN: We were founded in 2005 and we're a U.K. based company. We secure voice calls [and text messages] on cell phones. These are standard off-the-shelf smartphones, and we provide government-grade encryption on those voice calls to protect calls from the threat of interception. We have about 50 employees, and our main office is in London. Cellcrypt, Inc. is based in Washington and we also have offices in Dubai and Miami.

> **OK. Cellcrypt has developed secure voice encryption applications for Android-based smartphones, in addition to the iPhone and Blackberry. Could these be used by military troops in Afghanistan, for example?**

MEAKIN: Most of our business is government, military, law enforcement, though we also have a corporate division. So yes, we can and do support customers in the military. A couple of examples of this are welfare calling – where troops call back to their loved ones; another example is using the encryption application in covert operations by Special Forces. In covert operations, Cellcrypt smartphone encryption is particularly useful and fit for purpose because the voice encryption application can be downloaded over the air in minutes and used on a standard smartphone. So if you were in a covert operation, you could buy a standard smartphone from a store, connect to the Internet, and install the encryption application. It is inconspicuous when you use it – it just looks like a normal phone, rather than a military-style green device.

> **What is your primary motivation behind encrypting voice (and text messaging) in smartphones?**

MEAKIN: The main use case is covering "sensitive but unclassified" information.

This includes conversations with personnel, discussing plans – where to meet, who will be at the meeting, travel plans, ancillary information that becomes very important because [cell phone hackers] can then track people, track meetings, and that can become quite a powerful risk. And probably 95 percent of conversations in the government are of a sensitive but unclassified nature. Most likely all of those sensitive but unclassified conversations are made in the clear. Therefore, that is a market that Cellcrypt protects.

The reason we're doing this is because cell phones are vulnerable to interception, not just by traditional mechanisms of countries or well-funded criminals, but there's a new and emerging threat from hackers who have developed equipment that can intercept voice calls for as little as $1,500. It used to be hundreds of thousands of dollars – but now everyone can get ahold of it and so the risk of it happening has gone up.

> **Earlier you mentioned "government-grade" encryption. What kind of encryption do the Cellcrypt smartphone encryption applications use?**

MEAKIN: We use encryption certified by NIST: FIPS 140-2, again for sensitive but unclassified information. Our application uses NSA Suite B crypto. And we use two algorithms for every cryptographic process, For a voice call, we first use RC4 256-bit to encrypt; then we encrypt it again using AES 256-bit.

Something else very important about our approach: We don't have a central key server. Keys are managed directly between the recipient's and caller's applications. The key exchange is done by the software at both endpoints. There's no human interface into the keys. Therefore, there's a reduction in vulnerabilities because the key server can't be compromised because there is no key server.

❯ **Though all the Cellcrypt smartphone encryption apps secure voice calls, looks like Cellcrypt mobile encryption app for Blackberry is the only one providing security for text. Why is that?**

MEAKIN: The Blackberry application does have the additional feature of secure messaging. It's not an encrypted SMS system, and it's not an encrypted instant messaging system. It simply replaces the voice data with text data. Generally our system is used for perishable, [real-time] information.

The text is an alternative to voice – you might be typing in precise information such as an address, or text might be used as a backup to voice when a network in far-flung places makes latency on calls a couple seconds longer. Text is planned for other products as well at various stages in the future, including iPhone, Android, and Nokia smartphones, for which we also have a voice encryption application.

❯ **Wouldn't somebody just use encrypted SMS?**

MEAKIN: They could indeed, but ours is a different approach and designed to be used in different circumstances. For example, we avoid a "store and forward design." When you send an SMS, it uses store and forward technology: The message is sent to a central server that stores the text, then attempts to send it to the receiver. If the recipient is on a plane, they get it 3 hours later, after the plane lands, because the SMS was stored and forwarded. That infrastructure guarantees delivery. In our world, that represents a potential vulnerability because we're talking about perishable, real-time information: We don't necessarily want information stored and forwarded later. A good example is a covert operation when the target has moved and someone watching in a coffee shop sends a message "Target moved, fire now." For whatever reason, if the network isn't available, you don't want that message to get to the recipient half an hour later. There are operational consequences to that. Sometimes it's better for the message not to get through at all. But if people want encrypted SMS as a requirement, we work with partners to deliver encrypted SMS store and forward.

❯ **How does your system know when a text message can't be delivered, and how does your system cancel the message?**

MEAKIN: It knows that the recipient isn't online so it doesn't try to deliver it in the first place. It's like a voice call; for me to do that, I contact the recipient and do a cryptographic exchange and share a secret key. We have to both be online to do that. The secret key encrypts voice on my phone and decrypts it on the recipient's phone so he can hear it. The other thing is that we authenticate each other. It's the same with the text message. If the other person is not online, the text can't be sent in the first place.

> ❝ It's one of those untalked-about truths: that even though everyone wants to be secure, if they haven't got systems that are convenient, then they'll just use their cell phone. ❞

However, the user can set an automatic retry period, which is user defined. Otherwise, the retry time is 0. There might be circumstances where I say, "Now the person is going through a few tunnels and I'll retry it for the next 5 minutes or maybe an hour." I control the operational risk of the message not being delivered right now. If the recipient can't be reached in the retry period, the message evaporates.

❯ **That assumes you exchanged keys earlier?**

MEAKIN: Yes. When I first call someone using Cellcrypt and we've never contacted each other before, we exchange public keys. The public key I give to somebody else is what they use to encrypt the voice they send back to me, which my private key decrypts. It can only be decrypted by my private key. After a call has been made and we end the call, I have the option of retaining the other person's public key in my address book. [Thereafter], this can be used to exchange messages. When I send a message to him, I can already encrypt it with his public key. But I still need to connect with him for mutual authentication to make sure it's the right person

I'm talking to and derive a secret key to encrypt the text. And, as I mentioned, we both need to be online to do that.

❯ **The Cellcrypt data encryption only secures text, not emails or other data?**

MEAKIN: Yes, as I mentioned, it's not encrypted SMS, it's not encrypted instant messaging, and it's not encrypted Email. Each of those systems has its own architecture, which implies certain things. Those all have guaranteed delivery and store and forward architecture, which, again, we don't have. We are only sending alphanumeric strings in place of voice.

❯ **Will you be adding encryption for emails and other data to the Cellcrypt application?**

MEAKIN: No because those are provided by our partners. One thing our technology does really well is deliver information in real time. That's what the focus of our patents and innovation is. A voice call obviously is in real time, and you want the performance to be exceptionally good. So we've tuned our technology to deliver that. Email and SMS are not real time. It doesn't matter if they are delivered a minute later, so we don't need to tune that to real-time performance data. However, we are interested in transmitting other data – including video – [and enabling secure] file transfers in the future.

❯ **How does the Cellcrypt encryption application work?**

MEAKIN: The application that works on the smartphone itself is straightforward. Just click on the Cellcrypt icon to use it. It is a user interface that emulates a normal cell phone call and address book: You dial, speak, hang up, and so on. Its key strength is the crypto system behind the scenes. But the user experience is a very easy thing. That's important because if the system is not easy to use, people won't use it, even if their operational security policy says they must use it. It's one of those untalked-about truths: that even though everyone wants to be secure, if they haven't got systems that are convenient, then they'll just use their cell phone.

> **Which kinds of networks does the Cellcrypt secure smartphone application run on?**

MEAKIN: Effectively, this is an encrypted Voice over IP [VoIP] call – although we used our own protocol. So it will run on any network with an IP connection: GSM, CDMA, 2G, 3G, 4G, mesh networks, Wi-Fi, and satellite.

> **Is it possible that some parts of the world could not support that type of network?**

MEAKIN: Yeah somewhere, some places, but most countries have data cellular networks. There are very few places that don't. The biggest problem is the quality of the network. Some places in Africa and Latin America only have 2G networks and are bandwidth constrained. It gets a bit tricky. But one of the advantages of our product is we can operate at a very small bandwidth and over a very small 2G GPRS network. GPRS was the first data network applied to cellular. If there is no cellular coverage at all, some customers use satellite.

> **Tell us about the architecture of the Cellcrypt smartphone encryption application.**

MEAKIN: You have the application installed on both smartphones and a call is encrypted end-to-end between them. Both applications are connected to a signaling server in-between. When you make a call, one application asks the server if the other application is available to receive a call – as well as other things like authenticates the user and checks the license. If it is available, the server sets up the applications to talk directly to each other and then they perform the cryptographic handshake and make the call. The call is routed over our ECDN [Encrypted Content Delivery Network], which is a series of servers around the Internet that optimizes routing and call performance. Importantly, encryption/decryption is only done at the endpoints by the applications – the servers just set up and route calls and carry encrypted data.

We also have an endpoint application that sits on a server in the office and interfaces to a PBX. This allows an encrypted call into the office and can connect to landlines, voicemail, and conference calling facilities provided by the PBX.

> **What about latency and voice quality?**

MEAKIN: We are constantly using our innovation and technology to overcome the challenges. Our target is to have crystal-clear transmission with zero latency. That's a nice objective. It's not practical for that to happen all the time, just like any other communication mechanism. Whether it's cellular or a landline call, there are always trade-offs between latency and voice quality. But that's really what our core competency is. If I'm calling from London on a 3G network to Miami on a 3G network, we've recorded latency as 151 milliseconds. That's unnoticeable to somebody who's talking. The international standard for landline calls is about 300 milliseconds.

Because it's an IP-based call, the speech quality is very good because the packets get recombined at the recipient's end. But bandwidth changes and the inherent latency of the network changes constantly – which can affect the call. So a 2G GPRS call might have latency go down to 2 or 3 seconds in some circumstances. We intend to keep speech quality as high as possible and trade-off latency because we think it's more important to get an accurate voice message to the receiver than it is to get it there half a second sooner. If you're on Wi-Fi or 3G, latency is good for a normal voice call. It's only if you get to 2G GPRS or EDGE, that it might get out to 1, 2 seconds. It is a variable that is different for every person.

> **What's the future of the DoD in regard to ordering smartphones for troops in the field? Or are you more targeting benign DoD workers for your apps?**

MEAKIN: They will, they are doing. And as smartphones roll out in the DoD for data services, we will be looking to help provide a solution for securing the voice part.

But really, anybody who uses a smartphone needs a secure layer on it, if they've got sensitive information they want to protect. Talking about customers, if you were traveling to Latin America, you might need it so you don't get targeted for kidnapping. It's applicable really to anyone talking about something that can be exploited by someone else. It could be when, where, and with whom a meeting will be held, as I mentioned.

> **Will your encryption applications cover classified information, too, in the future?**

MEAKIN: In other regions, we do have governments using it for classified information. Those governments might have slightly different security requirements than someone like the U.S. or the Western/NATO governments. The security requirements of individual nations are specific to them. But I think it's worth looking at. The reason is cost of implementation. If you can deliver security over cell phones, you're solving a couple problems. One is you're using COTS devices, and they're cheap. The second thing is they're interoperable; everything works with everything else. If you provide a security layer, it makes it very easy.

I think that [some] governments are and will be looking at hardening the system such that classified information can be discussed. Obviously one of the downsides of using a COTS smartphone is there are other things to do to it, such as make the hardware more robust and make sure it becomes tamperproof. The downside is that it's open, and it can be tampered with. Certainly we need to lock that down. And plans can be put in place by governments to make sure that the overall security system in COTS smartphones is secure enough for higher classifications. It is an achievable project. **MES**

Ian Meakin is VP of Marketing at Cellcrypt. He has spent more than 20 years bringing high-tech solutions to governments and corporate enterprises for leading U.S. and U.K. companies. He has been in the mobile application industry for the past 7 years and holds a postgraduate Diploma in Marketing and BEng (Hons) in Aeronautical Engineering. He can be contacted at ian.meakin@cellcrypt.com.

SWAaP – Bringing Assured computing to SWaP with Freescale's QorIQ platforms

*By Glenn Beck (pictured),
Single Board Computer and
Aerospace/Defense Segment
Marketing Manager and
Geoffrey Waters, System
Architect, QorIQ Processors*

Size, Weight and Power (SWaP) have long been the primary parameters in designing electronic systems for military and aerospace applications. Previously, the major requirement to be balanced against SWaP was performance. That is no longer the case, as electronic systems of the future are increasingly required to be more open, upgradeable, exportable, and networked. These requirements create the need for another fundamental system design parameter: *Assurance*, making *SWAaP* the new paradigm in mil/aero electronics.

Now, systems must not only meet the performance of the mission, but must also protect against threats. Those threats come in three forms:

1) **Theft of functionality** – Loss of control of the system's functionality

2) **Theft of data** – Loss of data to an unauthorized party

3) **Theft of uniqueness** – Loss of product differentiation through reverse engineering, duplication, and unapproved interoperability

Freescale's QorIQ Communications Platforms, built on Power Architecture® technology, have intentionally implemented hardware and software features that provide protection against these threats. In doing so, they help decrease the cost and increase the security of systems through the elimination of Trusted Processors Modules (TPMs), FPGAs, and ASICS, which help lower cost and with fewer devices decrease exposed buses that are vulnerable to snooping.

Freescale's QorIQ platform's trust architecture allows for optimization of SWAaP, while delivering NSA recommended assurance features needed to truly leverage the performance, networking, and programmability features of multicore SoCs, while also providing a level of anti-tamper for intellectual property protection. The trust architecture in the QorIQ platforms provide protection through:

• **Secure boot** – Helps ensure that you start with a trusted code base or you don't start at all. This feature is designed to allow the user to implement an RSA public key to decrypt the signed hash while simultaneously recalculating the SHA-256 hash over the system code. The QorIQ processor compares the decrypted original hash to the freshly calculated hash – if the values match, the code is considered authentic. The user has provided and fused into the processor his public key, which cannot be read. By allowing the OEM the ability to fuse his public key into the processor, he has implemented the chain of trust.

• **Domain separation** – Allows the user to create partitions of cores, I/O, and memory regions to be separated from other partitions in the multicore processor. This operation occurs after the user has brought up the processor in a secure, trusted fashion. With the use of the hypervisor, the user determines the policies for each of the partitions that determines who has access to the assigned I/O and memory regions. This helps ensure that the operation of one partition does not interfere or corrupt another partition.

• **Threat protection** – When the processor is made aware of physical attacks to the system, it takes action to ensure the processor is unable to be utilized by an unauthorized third party. When a physical attack occurs to a system, a signal is sent to the processor at which time all of the memory regions in the device are written with zeros. Therefore, one can be assured that the processor is unusable. The device cannot be booted due to the secure boot feature. In addition, network attacks can be identified to prevent a third party from changing instruction code in the device. Each partition is given a logical numerical signature. If a packet does not have that identifier, it is not allowed into the domain. A security monitor is also reviewing activity in memory regions and if a policy is violated, it may choose to take a partition off-line per the policies implemented at the time of secure boot.

• **Secure debug** – Provides a question/answer sequence that is programmed into the boot fuse. This question/answer sequence is used to gain access through the debug port, allowing secure entry into the processor to debug systems or upgrade software remotely in the field.

The QorIQ communications platforms not only meet the Size, Weight, and Power requirements, but now bring Assured computing to these critical missions as well. SWAaP is now the standard that guides system architectures.

For more information, go to freescale.com/QorIQ.
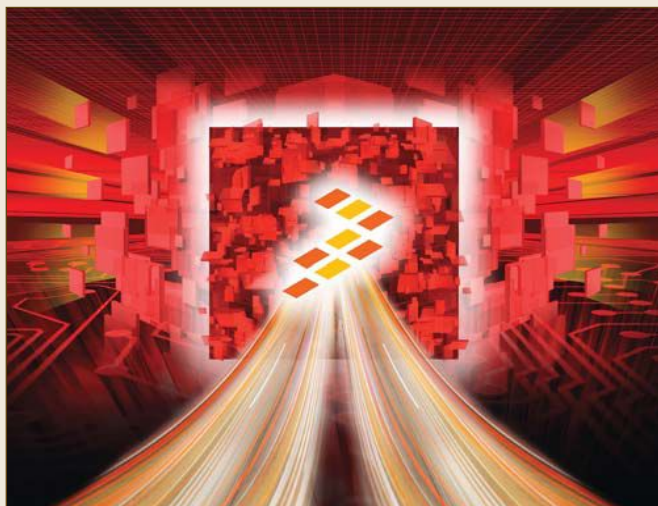
freescale™

Security including anti-tamper capabilities is critical for next-generation military systems

By John Costello, Vice President, Business Development, Military & Aerospace

There are many ways that military technology can be attacked or stolen. Attacks can be conducted remotely, and military technology also can be compromised following foreign sales to an ally, accidental loss, or capture by an enemy during a conflict. Because U.S. military hardware and software have a high technical content that provides a qualitative edge, protection of this technological superiority is a high priority, and security is a vital component in system design. The U.S. Department of Defense requires multiple layers of defense to ensure adequate protection of critical technology.

The security boundary is an important concept, especially for systems that, increasingly, must communicate with the outside world using the Internet and might not be located in a secure location. Indeed, a would-be attacker might have ready access to the hardware. FPGAs based on Static Random Access Memory (SRAM) are especially vulnerable to IP theft and product tampering. Because these devices cannot hold state when power is removed, a programming bitstream must be read into the device at startup. It is a simple matter to intercept this bitstream and clone the design, a process in which a pirate copies a board design, captures the configuration bitstream of the SRAM FPGA, and then builds an exact replica of the system without having to understand any of the details about the logic contained in the FPGA.

There are several ways to protect FPGAs from these attacks, including the incorporation of a cryptographic engine to decrypt the incoming bitstream. For even better security, it is necessary to include nonvolatile decryption key storage using antifuse or flash technology. Additionally, FPGAs and customizable System-on-Chip (cSoC) devices can be enhanced to include anti-tamper security to prevent or delay exploitation of essential or critical technologies. This security can also include Physical Unclonable Function (PUF) technology, which extracts a unique secret key from the cSoC and FPGA silicon hardware, rather than loading an externally generated key to on-chip non-volatile or battery-backed memory.

The memory function is another key area to consider. Military-grade solid-state disk drives protect sensitive data from environmental and human threats by incorporating a number of features including hardware-implemented AES-256 encryption with XTS, and hardware-based authentication. Other key requirements include: purging technology, which erases keys in less than 30 milliseconds to render data forensically unrecoverable; hardware-based fast clear of the entire drive within less than four seconds; and key management features that can be customized for encryption applications. Each of these features contributes to robust and highly effective tamper resistance.

Tamper-resistant technology is also the foundation for Information Assurance (IA), or the ability to secure data that is transmitted between systems. According to the 2007 ICTeR ANR Project report entitled, "SP1 Deliverable – State of the Art," there are five essential principles of IA that guarantee the correct execution of a program and the correct management of the communications:

- **Confidentiality:** Only the entities involved in the execution or the communication can have access to the data;
- **Integrity:** The message must not be damaged during transfer or the program must not be altered before its execution;
- **Availability:** The message or program must be available;
- **Authenticity:** The entity must be sure that the message comes from the right entity or the system must trust the program source code; and
- **Non-repudiation:** The entities involved in the execution or communication must not have the possibility to deny the exchange.

Effective IA requires a combination of approaches, including software-based key security, which protects private repository keys, public authentication keys, and the key-based authentication techniques used to bind software to a unique hardware identifier. This can be accomplished by transforming the decryption key into a series of algebraic keys. Concealing these keys is critical to preventing system compromise from outside forces.

Another key IA layer is hardware protection. Tamper-resistant hardware provides a secure computation environment. The security boundary may enclose a single integrated circuit, such as a flash-based FPGA or microcontroller, or it may comprise multiple chips in a secured module. Tamper-resistant features prevent reverse engineering by providing static and dynamic tamper prevention, detection, and reaction using both passive and active mechanisms. The hardware not only provides the secure platform to store and manipulate data using secret keys, but also provides additional protection against vulnerabilities such as, but not limited to, side-channel attacks such as those based on differential power analysis.

A third IA requirement is software protection, which combines strong encryption, anti-debugging, anti-tampering, and real-time execution monitoring. This enables software applications to defend themselves and take appropriate action to protect against reverse engineering, code lifting, and unauthorized changes, with minimal impact on runtime performance.

The ability to more effectively, efficiently, and securely move information through the military operation is critical to maintaining a war-fighting advantage. Security and tamper resistance are critical in applications including aircraft and aerospace; missiles and ordnance; and Command, Control, Communications, Computers, Intelligence, Surveillance, and Reconnaissance (C4ISR). Equipment can be protected at the FPGA and memory level, and also system-wide. This comprehensive, layered approach is necessary to effectively protect modern military electronics from a variety of increasingly critical threats.

Microsemi.

EXECUTIVE SPEAKOUT

Military EMBEDDED SYSTEMS

RESOURCE GUIDE PROFILE INDEX

Pinnacle Data Systems, Incorporated

6600 Port Road • Groveport, OH 43125 USA
Tel: +1 (614) 748-1150 • Fax: +1 (614) 748-1209
www.pinnacle.com

ATCA-F1 Dual Socket AMD® Opteron Blade

This industry-proven ATCA-F1 server blade features two AMD Opteron™ dual-, quad-, or six-core processors for significant performance available in a general-purpose AdvancedTCA® compute board. Now offered with a full 32GB of memory and the benchmark-setting AMD six-core CPUs, this blade supplies the extra level of computing horsepower and built-in virtualization support demanded by the next generation of COTS integrated architectures. With its robust design, the ATCA-F1 blade has been thoroughly tested and successfully deployed in critical systems, including military applications.

The ATCA-F1 features a standard Zone 3 interface for connection to PDSi's ATCA-RT01 Rear Transition Module (RTM), which adds SAS storage, video, Ethernet and USB resources. The ATCA-F1 and ATCA-RT01 RTM combination has been validated and is hardware-compliant with the VMware® ESX Server virtualization platform.

FEATURES

> AMD Socket F (1207-pin) CPU sockets
> Dual-core, quad-core, or six-core AMD Opteron CPUs
> Zone 3 RTM interface (to ATCA-RT01 RTM)
> 4 x DIMM sockets, up to 32GB DDR2, 667 MHz Memory
> Ethernet Base and Fabric Interfaces
> 2 x Gb Ethernet links (front panel)
> 2 x USB 2.0 (front panel)
> 1 x AMC.1 slot (mid-size)
> VMware-compliant
> Pigeon Point IPMC management
> Customization and 3rd party integration welcomed

For more information, contact: info.sales@pinnacle.com

www.mil-embedded.com/p46014

Pinnacle Data Systems, Incorporated

6600 Port Road • Groveport, OH 43125 USA
Tel: +1 (614) 748-1150 • Fax: +1 (614) 748-1209
www.pinnacle.com

Dual Intel® Sandy Bridge ATCA® Blade (ATCA-N1)

PDSi's Dual Intel Sandy Bridge AdvancedTCA® (ATCA) Blade with RTM Interface (ATCA-N1) is a high-performance computing platform for ATCA systems. Designed to provide maximum computing, memory and storage flexibility in a single ATCA slot, the ATCA-N1 features dual CPU sockets that can be populated with Intel Sandy Bridge CPUs, supports up to 192GB of DDR3 1333MHz memory and can be configured into an enterprise configuration (see ATCA-N2) for higher level computing requirements.

Other features include support for up to four SATADIMM solid state drives, 400W per ATCA slot chassis implementations and a Zone 3 interface for connection to a Rear Transition Module.

PDSi can also provide customization, ruggedization, turnkey integration and support for ATCA systems, as well as extended warranty and repair services.

FEATURES

> Intel-based AdvancedTCA blade server utilizing Intel's 32nm processor technology
> 10Gb Base and Fabric interfaces
> Pigeon Point IPMC management
> 2 x Sandy Bridge CPU sockets
> 12 x DIMM sockets, up to 192GB DDR3 1333MHz (4 x dual purpose)
> Supports up to 4 x SATADIMM SSD, MLC and SLC versions
> Front Panel I/O: 1 x 1Gb Ethernet, 4 x USB 2.0, 1 x Serial
> Zone 3 interface

For more information, contact: info.sales@pinnacle.com

www.mil-embedded.com/p53128

Kontron

14118 Stowe Drive • Poway, CA 92064-7147 USA
888-294-4558
www.kontron.com

Kontron AdvancedMC™ processor module AM4120

The Kontron AdvancedMC™ processor module AM4120 is equipped with the Freescale™ QorIQ P2020 dual core processor with up to 1.2 GHz and based on the Power Architecture® e500. Application reliability is ensured with features such as the redundant universal bootloader U-Boot and persistent memory for cycle data storage. A dedicated Module Management Controller (MMC), which supports basic IPMI commands, is used for board management, enabling operators to monitor the status of the AdvancedMC™ module in the system, simplifying system management and improving availability. The Kontron AM4120 also offers increased longevity due to Freescale's processor availability until at least 2018, careful component selection and a microSDHC card socket that is not impacted by regular Flash discontinuation.

The universal Kontron AM4120 offers cost-sensitive AdvancedTCA® and MicroTCA™ system designs, long-term availability, support of various fabrics and an extraordinary performance-per-watt ratio.

FEATURES

› Freescale™ QorIQ P2020 – dual core @ 1.2 GHz
› Up to 4 GB RAM with ECC
› 3x GbE: AMC port 0 + 2x front, or port 0, 1 + 1x front
› microSD card socket
› Extended temperature range on project request
› Useable in managed (with MCH) and unmanaged (without MCH) systems
› Processor roadmap until 2018

For more information, contact: info@us.kontron.com

www.mil-embedded.com/p52841

RadiSys Corporation

5445 N.E. Dawson Creek Drive • Hillsboro, OR 97124 USA
503-615-1100 • Fax 503-615-1115 • Toll-Free 800-950-0044
www.radisys.com

RadiSys COM Express CEQM57XT

Next-Generation Performance

The RadiSys COM Express CEQM57XT combines the next-generation performance Intel® Core™ i5 and i7 processors and the Mobile Intel® QM57 Express Chipset with RadiSys -25C to +70C extended temperature and vibration specifications to provide breakthrough processing performance on a ruggedized COM Express module. The basic size 95mm x 125mm module is ideal for compute-intensive applications such as military-aerospace, transportation and industrial applications requiring high levels of performance and durability.

FEATURES

› -25C to +70C extended temperature range
› Intel® Core™ processor options:
  • Intel Core™ i7 610E 2.53GHz
  • Intel Core™ i7 620L 2.0GHz
  • Intel Core™ i7 620U 1.06GHz
› Mobile Intel® QM57 Express Chipset
› Dual-channel DDR3, up to 8GB
› Type 2 pin-out
› TPM
› Six PCI Express x1 ports, one PCI Express x8 port
› Single or dual Gigabit Ethernet options
› Conformal coat option – Humiseal 1B31

For more information, contact: info@radisys.com

www.mil-embedded.com/p47710

Pinnacle Data Systems, Incorporated
6600 Port Road • Groveport, OH 43125 USA
Tel: +1 (614) 748-1150 • Fax: +1 (614) 748-1209
www.pinnacle.com

AMD® Socket S1 COM Express Module

PDSi's AMD Socket S1 COM Express Module (COMX-S1) is a low-cost, compact, embeddable computing core with the capability to drive a broad range of OEM applications. Built around AMD's x86-based Socket S1 processors, it enables 64-bit computing at a progression of performance levels from the ultra low-power AMD Sempron™ 2100+ (perfect for fanless applications) to the dual-core muscle of the AMD Turion™ X2 TL-62.

The Computer-on-Module design is fully compliant with PICMG COM Express Type 2 specifications, offering a large variety of interfaces to cover the needs of most embedded applications. The AMD M690/SB600 chipset provides I/O versatility, with interfaces for high-speed data, audio, high-resolution graphics, and more. PDSi's companion COM Express Carrier Board (COMX-CB) provides a plug-and-go platform with physical ports for virtually all of the I/O types supported by the COM Express Module.

FEATURES
> Supports AMD Socket S1 processor family including AMD Turion™ 64 X2 dual-core mobile and Mobile AMD Sempron™
> Up to 4 GB DDR2 SDRAM
> Video – Dual-channel LVDS, VGA, optional TV Out
> AC'97 audio interface, optional HD Audio
> 4 x PCI Express interfaces (x8, x1, x1, x1)
> 8 x USB ports
> 4 x SATA II ports
> 1 x Ethernet port (10/100/1000)
> PICMG COM Express Type 2-compatible pinout
> MicroATX Carrier Board available for rapid startup
> Extended availability assured; 3rd party integration welcomed

For more information, contact: info.sales@pinnacle.com   **www.mil-embedded.com/p35178**

ACTIS Computer
1898 E. Southern Ave. • Tempe, AZ 85282 USA
Tel: 480-838-1799 • Fax: 480-838-4477
www.actis-computer.com

6U CompactPCI Single Board Data Recorder CSBX-6315

ACTIS Computer announces the new CSBX-6315: a data Recorder on a standard 6U cPCI single board, built around the third-generation PowerQUICC 2 CPU, designed by Freescale Semiconductor. This new CSBX-6315 supports processor speeds up to 400MHz, allowing a trade-off between lower power and higher performance.

Block Diagram Overview
ACTIS engineers have designed the CSBX-6315, a versatile 6U cPCI Single Board Computer for critical embedded applications with a power dissipation of less than 8 watts.

Onboard memory includes 256 Mbytes of soldered DDR SDRAM and 8 Mbytes of NOR flash memory for the onboard monitor software. In addition, a 1 GByte NAND flash is provided for additional data or program storage.

FEATURES
> 6U Single-slot cPCI
> Freescale third generation PowerQUICC 2
> Gigabit Ethernet port (on the front panel)
> 3 Fast Ethernet ports
> 2 serial ports
> 3 High-speed USB 2.0 ports
> 2 SSD Slots: up to 64 GBytes on SATA and mSATA
> Expansion Slots: one PCI Express Mini and one PMC Slot

For more information, contact: sales-usa@actis-computer.com   **www.mil-embedded.com/p53160**

CelDel, LLC

1 Jonathan Bourne Drive, Suite 8 • Pocasset, MA 02559 USA
508-718-2050
www.celdel.com

Where hardware & software overlap

Freescale Dual P2020 3U CPCI Board

The Dual P2020 3U CPCI Board delivers very high processing density in a single, compact and rugged 3U CPCI form factor. This powerful platform is ideal for complex embedded applications such as secure networking or data processing.

CelDel provides full Linux with multiprocessor support, delivering a powerful embedded software development environment for rapid porting of custom applications. Be up and running immediately, utilizing the existing HTTP, FTP, NFS, Samba, DHCP and Telnet servers.

The Dual P2020 3U CPCI Board is a ruggedized, ROHS-compliant, hot-swappable CPCI form factor and provides flexibility and scalability for the most demanding systems.

About CelDel:

CelDel LLC is a leading supplier of standards-based embedded hardware and software. Our team serves the medical, military, commercial, instrumental and consumer electronics industries.

FEATURES

› Two Dual-Core P2020 processors running up to 1.2GHz
› Four 1000BASE-T ports via front panel with support for up to eight 1000-BASE-T ports (rear I/O support)
› Two USB2.0 host ports via rear transition I/O
› 128MByte (x16) NOR boot flash
› 4GB DDR2-667MHz SODIMM with ECC (x72) – 2GB per processor
› Up to two SATA interfaces via rear transition I/O

For more information, contact: octavio.amaral@celdel.com

www.mil-embedded.com/p46396

COTS Collection: Boards, Carriers, and Mezzanine: CompactPCI

Emerson Network Power

2900 S. Diablo Way, Suite 190 • Tempe, AZ 85282 USA
1-800-759-1107 or +1-602-438-5720
www.Emerson.com/EmbeddedComputing

CPCI7203 Air- and Conduction-Cooled 3U Processor Board

The CPCI7203 3U SBC features the integrated dual-core Intel® Core™ i7 processor for use in high-performance, space-constrained applications. This leading-edge thermal and rugged solution makes it ideal for harsh environments. On-board memory includes up to 8GB DDR3 and additionally on the air cooled variants, 256KB non-volatile F-RAM and 4GB MicroSD flash.

Connectivity is optimized for maximum throughput and flexibility. The air-cooled variants have two Gigabit Ethernet ports, two USB 2.0 ports, and one VGA on the front panel. Rear I/O includes one serial port, two SATA ports and four PCI interfaces to the rear. Conduction-cooled variants also provide rear VGA. The Trusted Platform Module (TPM) enhances data security and encryption capabilities.

This low-power, high-performance SBC offers full hot swap compliance per PICMG® 2.1; supports the PICMG 2.9 System Management spec; & is pin compatible with PICMG 2.30 CompactPCI PlusIO spec. It supports a range of operating system & software options.

FEATURES

› Integrated dual-core Intel® Core™ i7 processor (up to 2 GHz)
› Up to 8GB ECC-protected DDR3-800/1066 (soldered)
› 256KB non-volatile F-RAM (air cooled)
› Mobile Intel® 5 Series chipset: Ibex Peak-M PCH
› One VGA, two USB 2.0 & on-board Gigabit Ethernet interfaces
› 4GB MicroSD
› Air- and conduction-cooled
› Extended temperature range (-40 °C to +85 °C)
› Ideal for industrial, medical and military/aerospace applications (railway control, factory automation, semiconductor processing, robotics, image processing, vetronics, VoIP and first responder)

For more information, contact: EmbeddedComputingSales@Emerson.com

www.mil-embedded.com/p45456

Kontron

14118 Stowe Drive • Poway, CA 92064-7147 USA
888-294-4558
www.kontron.com

Kontron 3U CompactPCI CP3002 and CP3923

Processor board & Gigabit Ethernet switch

The Kontron CP3002-RC with the Intel® Core™ i7 mobile processor (620LE – 25 W TDP) is conduction cooled and supports operational temperatures ranging from -40° to +85°C according to VITA 47. The powerful processor and the ruggedness make the Kontron CP3002 ideal for a range of new real-time, data-intensive and network-centric applications in the defense and aerospace industry as well as in transportation and other sectors that require reliable performance in harsh environments. Up to 8 GB of soldered DDR3 1066 MHz ECC memory ensures data accuracy for demanding and mission-critical and safety-critical applications.

The Kontron CP3923 is a fully managed layer 2/3 Gigabit Ethernet switch offering IPv4/IPv6 routing and full management capabilities. It supports a powerful set of CLI, Telnet, web and SNMP management interfaces to configure the entire set of protocols and parameters including Layer 2 and Layer 3 (IPv4/IPv6) protocols, multicasting, QoS and security.

FEATURES

> **CP3002-RC**
   • Rugged, conduction cooled
   • Operating temperature -40° to +85°C (VITA 47)
   • 8 GB soldered memory ECC at 1066 MHz
   • Comprehensive feature set – 4x GbE, 2x USB 2.0, VGA, 2x COM

> **CP3923**
   • Fully managed Layer 2 and Layer 3 switching and routing
   • Leading-edge technology based on BCM56226
   • Versatile design with RJ-45 or M12D front options
   • EN50155 compliant (with M12 versions)
   • Design prepared for conduction cooling

For more information, contact: info@us.kontron.com   **www.mil-embedded.com/p45804**

Kontron

14118 Stowe Drive • Poway, CA 92064-7147
888-294-4558
www.kontron.com

Kontron 6U CompactPCI CP6003-SA and CP6930

Processor Board & 10 Gigabit Ethernet Switch

The Kontron CP6003-SA, available with the 2nd Generation Intel® Core™ i7/i5 quad and dual core processors, meets the highest computing and graphics performance requirements at the lowest power budgets (32 nm technology). The power of up to four cores enables virtualization and multithreading applications to run in full 64-bit mode using Intel® Enhanced Virtualization, HD Graphics and Turbo Boost Technology.

The Kontron CP6930 is a hot-swappable CPCI and VITA 31 switch with 26 GbE ports and six high-performance uplinks (10 GbE). The CP6930 is designed for future-oriented applications requiring out-standing bandwidth and communication safety. Six SFP+ front ports running at 10 Gb/s full line speed and two SFP ports at the front providing 1 Gb support give flexibility and an intrinsic value for today's investments.

FEATURES

> **CP6003-SA**
   • Up to 16 GB DDR3 1600 MHz memory with ECC via two SO-DIMM sockets
   • 5x Gigabit Ethernet interfaces via PCI Express
   • Latest I/O technology with SATA Gen 3, RAID, SSD
   • XMC/PMC + onboard HDD + SSD flash configuration possible

> **CP6930**
   • Non-blocking Layer 2 and 3 switching and routing, IPv6-compliant
   • 24x Gigabit Ethernet ports: PICMG 2.16/VITA 31.1-accordance
   • Fully managed, hot swap, IPMI
   • DHCP server and boot media support
   • Comprehensive firmware package

For more information, contact: info@us.kontron.com   **www.mil-embedded.com/p53035**

Pinnacle Data Systems, Incorporated

6600 Port Road • Groveport, OH 43125 USA
Tel: +1 (614) 748-1150 • Fax: +1 (614) 748-1209
www.pinnacle.com

CP86-N1 Intel® Core™ 2 Duo Processor Blade

PDSi's Intel-based CompactPCI x86 Processor Blade (CP86-N1) provides a robust, high-performance general purpose compute platform for use in CompactPCI PICMG 2.16 systems. This latest addition to PDSi's ComputeNode™ family of carrier-grade CompactPCI solutions is built around Intel's 45nm technology "Penryn" Core 2 Duo processor and server-grade "Eagle Lake" chipset (5100 MCH/ICH9R) supporting ECC memory. This powerful, compact blade offers the highest performance and dependability in its class.

The ComputeNode CP86-N1 blade includes a standard PMC/XMC site for I/O expansion and features an onboard SATA drive plus high resolution graphics. I/O capability covers a very broad range of interfaces that can be accessed through one of PDSi's companion Rear Transition Modules such as CP86-RT01. Two 1000Base-T Ethernet ports provide the PICMG 2.16-compliant fabric interfaces, making the CP86-N1 fully compatible with any ComputeNode cPSB chassis.

FEATURES

› Server-grade CompactPSB compute blade
› Intel T9400 Core 2 Duo 2.53GHz with 5100 MCH/ICH9R
› Up to 8GB Registered ECC DDR2 667 Memory
› 1 PMC/XMC mezzanine site standard
› 2 x 1Gb Ethernet links (front panel)
› 2 x USB 2.0 ports (front panel)
› Rear I/O interfaces
› Customization and 3rd party integration welcomed; extended availability assured

For more information, contact: info.sales@pinnacle.com

www.mil-embedded.com/p40980

ACCES I/O Products, Inc.

10623 Roselle Street • San Diego, CA 92121 USA
858-550-9559
www.accesio.com

Multifunction DAQ-PACK Series (Up to 128 Channels)

The DAQ-PACK Series is a highly integrated multifunction data acquisition and control system. The system offers an ideal solution for adding portable, easy-to-install high-speed analog and digital I/O capabilities to any PC or embedded system with a USB port. It performs signal conditioning such as RC filtering, current inputs, RTD measurement, bridge completion, thermocouple break detection, voltage dividers, small signal inputs, and sensor excitation voltage supply.

The small, compact, multifunction I/O DAQ-PACK Series units provide the user with everything needed to start acquiring, measuring, analyzing, and monitoring in a variety of applications. These data acquisition and control devices can be used in many current real-world applications such as precision measurement, analysis, monitoring, and control in countless embedded applications.

FEATURES

› 32, 64, 96, or 128-channel single-ended or differential analog inputs
› High-speed USB 2.0 multifunction DAQ
› Sustained sampling speeds up to 500KHz
› 12 or 16-bit resolution A/D converter
› Flexible, software-configured functionality
› 18 input ranges, 9 unipolar and 9 bipolar, per 8-channel programmable
› Wide range of flexible signal conditioning types
› Autocalibration and oversampling for real-time accurate data
› A/D starts via software, timer, or external trigger
› Two 16-bit analog outputs and 16 high-current digital I/O lines
› 16-bit programmable counter/timer

For more information, contact: contactus@accesio.com

www.mil-embedded.com/p52224

Annapolis Micro Systems, Inc.

190 Admiral Cochrane Drive, Suite 130 • Annapolis, MD 21401 USA
410-841-2514
www.annapmicro.com

Dual 4.0 GSps DAC

The Annapolis Micro Systems Dual Channel 4.0 GSps D/A I/O Card provides one or two 12-bit digital output streams at up to 4.0 GSps. The board has one or two Max 19693 for 4.0 GSps, Max 19692 for 2.3 GSps, or Max 5859 for 1.5 GSps.

The Dual Channel DAC board has five SMA front connectors: two single-ended DAC outputs, a high-precision trigger input with Fs precision, and a universal single- or double-ended 50-ohm clock input. It has excellent gain flatness in the first 3 Nyquist Zones, ultra-low skew and jitter saw-based clock distributions, and main board PCLK sourcing capability.

In concert with the WILDSTAR 4 or WILDSTAR 5 FPGA processing main boards, this mezzanine supplies user-configurable real-time A to D conversion and digital output. Up to two A/D or D/A and up to two serial I/O cards can reside on each WILDSTAR 4 or WILDSTAR 5 VME/VXS or IBM Blade main board, or up to one A/D or D/A and up to one serial I/O card on each PCI-X or PCIe main board.

Our boards run on many different operating systems. We support our boards with a standardized set of drivers, APIs, and VHDL simulation models. VHDL source is provided for the interfaces to A/Ds, D/As, DRAM/SRAM, LAD bus, I/O bus, and PPC Flash. CoreFire™ users will have the usual CoreFire Board Support Package.

The combination of our COTS hardware and our CoreFire FPGA Application Development tool allows our customers to make massive improvements in processing speed, while achieving significant savings in size, weight, power, person-hours, dollars, and calendar time to deployment.

Annapolis Micro Systems, Inc. is a world leader in high-performance, COTS FPGA-based boards and processing for RADAR, SONAR, SIGINT, ELINT, DSP, FFTs, communications, Software-Defined Radio, encryption, image processing, prototyping, text processing, and other processing-intensive applications.

Annapolis is famous for the high quality of our products and for our unparalleled dedication to ensuring that the customers' applications succeed. We offer training and exceptional special application development support, as well as more conventional customer support.

FEATURES

> One or two 12-bit Analog to Digital Converters: Max 19693 for 4.0 GSps, Max 19692 for 2.3 GSps, or Max 5859 for 1.5 GSps

> Five SMA front panel connectors: two single-ended DAC outputs, one high-precision trigger input with Fs precision

> One universal single- or double-ended 50-ohm clock input

> High-precision trigger input manufacturing options – 1.65 V LVPECL, 2.5 V LVPECL, 3.3 V LVPECL

> I/O card plugs onto WILDSTAR 4 or 5 VME/VXS/PCI-X/PCIe/ IBM Blade main boards

> JTAG, ChipScope, and Serial Port access

> Full CoreFire Board Support Package for fast, easy application development

> VHDL model, including source code for board-level interfaces

> Proactive thermal management system

> Industrial temperature range

> Includes one-year hardware warranty, software updates, and customer support

> Designed and manufactured in the USA

For more information, contact: wfinfo@annapmicro.com

www.mil-embedded.com/p36023

Military Embedded Systems   Resource Guide

Annapolis Micro Systems, Inc.
190 Admiral Cochrane Drive, Suite 130 • Annapolis, MD 21401 USA
410-841-2514
www.annapmicro.com

Dual 40/100G CFP IBM Blade 2nd-Slot Card

Annapolis Micro Systems, Inc. is releasing the first COTS board in the world capable of capturing and processing multiple 100Gbit signals in real time! The Annapolis Dual 40/100G CFP IBM Blade 2nd-Slot Card enables the capture and real-time processing of massive amounts of data for Network Security and Signal Intelligence applications. Annapolis has integrated two CFP interfaces into their Dual CFP IBM Blade 2nd-Slot Card, allowing the capture, buffering and processing or transmission of two 100Gbit Ethernet streams per card or four 40Gbit Ethernet streams per card. The Dual CFP IBM Blade 2nd-Slot Card features one or two CFP transceiver cage interfaces, which can support either 100Gbit or 40Gbit Ethernet, 100Gbit OTU4 or 40Gbit SDH/OTU3.

Integral to any massively bandwidth intensive application is the ability to process that data. The Dual CFP IBM Blade 2nd-Slot Card is designed to interface to the 8 processors available on the WILDSTAR 5 IBM Blade main board. Using Altera Stratix IV FPGA processors, up to 37 GBytes of DDR3 DRAM and a 160x160 6.5Gbit crossbar, the Dual CFP card is able to buffer and distribute the data from the CFP interfaces to the processors on a WILDSTAR 5 IBM Blade main board at full bandwidth.

Annapolis will supply the Dual CFP IBM Blade 2nd-Slot Card with Ethernet MACs for 40Gbit and 100Gbit Ethernet, and full standardized board support for VHDL models, drivers and APIs.
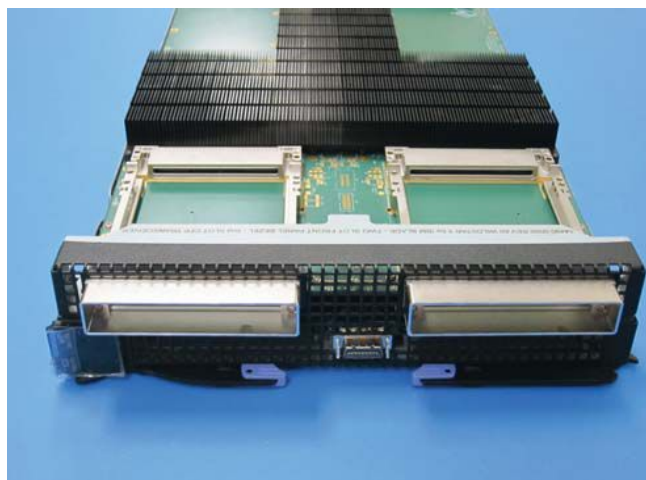
FEATURES

> One or two CFP interfaces, each supporting full duplex:
  – **100Gbit and 40Gbit Ethernet (802.3ba)**
  – **100Gbit (OTU4) and 40Gbit (SDH/OTU3) for Telecommunications**
  – **Three 40Gbit Ethernet streams**
> One or two Altera Stratix IV GT EP4S100G5 FPGAs
> Up to 37 GBytes of DDR3 DRAM arranged in up to eight 72-bit ports
> Integrated as 2nd slot of WILDSTAR 5 for IBM Blade main processor board
> Up to 68 Full Duplex Serial I/O lanes from CFP 2nd slot board to WILDSTAR 5 IBM Blade main board
> Supports enough memory bandwidth for buffering the incoming data
> Includes 100Gbit and 40Gbit Ethernet MACs for Altera Stratix IV

For more information, contact: wfinfo@annapmicro.com   **www.mil-embedded.com/p46786**

COTS Collection: Boards, Carriers, and Mezzanines: FPGA/Reconfig. Computing

Military Embedded Systems   Resource Guide

Annapolis Micro Systems, Inc.

190 Admiral Cochrane Drive, Suite 130 • Annapolis, MD 21401 USA
410-841-2514
www.annapmicro.com

2.0 GSps 10-bit A/D

The Annapolis Single Channel 2.0 GSps A/D I/O Card provides one 2.0 GHz A/D input with a resolution of 10 bits. The board has one e2v AT84AS004 that is fed by an onboard analog input circuit, which converts the single-ended 50-ohm SMA input into differential signals for the ADC. There is a universal single-ended 50-ohm SMA clock input and a high-precision trigger input allowing multiple A/D I/O cards to be synchronized together. Synchronization of A/D I/O cards can be facilitated by the Annapolis 4 or 8 Channel Clock Distribution Boards.

In concert with the WILDSTAR 4 or WILDSTAR 5 FPGA processing main boards, this mezzanine supplies user-configurable real-time continuous sustained processing of the full data stream. Up to two A/D and up to two serial I/O cards can reside on each WILDSTAR 4 or WILDSTAR 5 VME/VXS or IBM Blade main board or up to one A/D and up to one serial I/O card on each PCI-X or PCIe main board.

Our boards run on many different operating systems. We support our boards with a standardized set of drivers, APIs, and VHDL simulation models. VHDL source is provided for the interfaces to A/Ds, D/As, DRAM/SRAM, LAD bus, I/O bus, and PPC Flash. CoreFire™ users will have the usual CoreFire Board Support Package.

The combination of our COTS hardware and our CoreFire FPGA Application Development tool allows our customers to make massive improvements in processing speed, while achieving significant savings in size, weight, power, person-hours, dollars, and calendar time to deployment.

Annapolis Micro Systems, Inc. is a world leader in high-performance, COTS FPGA-based boards and processing for RADAR, SONAR, SIGINT, ELINT, DSP, FFTs, communications, Software-Defined Radio, encryption, image processing, prototyping, text processing, and other processing-intensive applications.

Annapolis is famous for the high quality of our products and for our unparalleled dedication to ensuring that the customers' applications succeed.

FEATURES

> One e2v AT84AS004 (2.0 GHz, 10-bit) A/D
> Four SMA front panel connectors: one 50-ohm analog input, one single-ended 50-ohm clock input, or differential 1.65 V LVPECL clock input
> One high-precision trigger input with Fs precision; high-precision trigger input – 1.65 V LVPECL, 2.5 V LVPECL, 3.3 V LVPECL
> Analog input bandwidth is 100 KHz-3.0 GHz
> I/O card plugs onto WILDSTAR 4 or 5 VME/VXS/PCI-X/PCIe/ IBM Blade main boards
> JTAG, ChipScope, and Serial Port access
> Full CoreFire Board Support Package for fast, easy application development
> VHDL model, including source code for board-level interfaces
> Proactive thermal management system
> Includes one-year hardware warranty, software updates, and customer support
> We offer training and exceptional special application development support, as well as more conventional customer support
> Designed and manufactured in the USA

Annapolis Micro Systems, Inc.

190 Admiral Cochrane Drive, Suite 130 • Annapolis, MD 21401 USA
410-841-2514
www.annapmicro.com

Four Channel Clock Synchronization Board

The Four Channel Clock Distribution Board distributes a common clock and synchronized control signal triggers to multiple cards in the system. This 6U VME64x/VXS board provides four high-speed, ultra-low jitter, ultra-low skew differential bulkhead mounted clock outputs, two ultra-low skew differential vertical SMA onboard clock outputs, and four ultra-low skew and clock synchronized single-ended bulkhead mounted control signal triggers.

A jumper set at board installation time or via optional P2 serial port determines which one of the two installed clock sources is active. Manufacturing options for Clock Source 0 are single-ended or differential external clock, a PLL ranging from 700 MHz to 3 GHz with an onboard reference oscillator, or a PLL ranging from 700 MHz to 3 GHz with a 10 MHz external reference. Manufacturing options for Clock Source 1 are a PLL ranging from 700 MHz to 3 GHz with an onboard reference oscillator, a PLL ranging from 700 MHz to 3 GHz with a 10 MHz external reference, or an onboard low frequency oscillator ranging up to 800 MHz.

The four control trigger outputs can originate from a high-precision external source via front panel SMA, from a manual push button on the front panel, or from software via an optional Backplane P2 Connector Serial Port. These trigger outputs are synchronized to the distributed clock to provide precise output timing relationships.

Annapolis Micro Systems is a world leader in high-performance, COTS FPGA-based boards and processing for RADAR, SONAR, SIGINT, ELINT, DSP, FFTs, communications, Software-Defined Radio, encryption, image processing, prototyping, text processing, and other processing-intensive applications.

Annapolis is famous for the high quality of our products and for our unparalleled dedication to ensuring that the customers' applications succeed. We offer training and exceptional special application development support, as well as more conventional support.

FEATURES

› Four synchronized differential front panel clock outputs up to 3 GHz with typical skew of 5 ps

› Ultra-low clock jitter and phase noise – 275 fs with 1,280 MHz PLL and external 10 MHz reference

› Onboard PLL's manufacturing options provide fixed frequencies of 700 MHz to 3 GHz, locked to internal or external reference

› Onboard low frequency oscillator provides fixed frequencies up to approximately 800 MHz

› Four synchronized trigger outputs, always synchronized with the output clock, with typical skew of 5 ps

› Jumper selectable trigger output levels of 3.3 V PECL, 2.5 V PECL, or 1.65 V PECL

› Source trigger from front panel SMA, push button, or optional P2 serial port

› Cascade boards to provide up to 16 sets of outputs

› Compatible with standard VME64x and VXS 6U backplanes

› Universal clock input supports wide range of signal options, including signal generator sine wave

› Differential clock input permits multiple standards including: LVDS, 3.3 V PECL, 2.5 V PECL, and 1.65 V PECL

› Clock and trigger outputs compatible with all Annapolis Micro Systems, Inc. WILDSTAR™ 2 PRO I/O Cards and WILDSTAR™ 4/5 Mezzanine Cards

COTS Collection: Boards, Carriers, and Mezzanines: FPGA/Reconfig. Computing

Military Embedded Systems   Resource Guide

Annapolis Micro Systems, Inc.

190 Admiral Cochrane Drive, Suite 130 • Annapolis, MD 21401 USA
410-841-2514
www.annapmicro.com

SFPDP UNI6 I/O

Annapolis Micro Systems Inc.'s FPGA-based WILDSTAR family provides 24 SFPDP channels per VME slot.

The Annapolis SFPDP cards (UNI3 or UNI6) come with an easy to use Serial FPDP interface supporting up to 12 lanes of 2.5 Gb full duplex data. Three frame types are supported: Normal Data Fiber Frame, Sync Without Data Fiber Frame, and Sync with Data Fiber Frame in Point-to-Point Mode.

The card has three individually configurable, industry-standard 4X connectors, providing four lanes per connector, with dedicated signal conditioners to ensure clean communication. It supports up to 7.5 GB full duplex per I/O card and a wide variety of readily available copper and fiber cables.

Up to two serial I/O cards and two LVDS I/O cards can reside on each WILDSTAR 4 or WILDSTAR 5 VME/VXS main board, with half that number for the PCI-X or PCIe. The SFPDP card (UNI6) supports RocketIO protocol at up to 75 Gb full duplex per I/O card, three ports of 10 G full duplex InfiniBand per I/O card, or 10 G full duplex Ethernet per I/O card.

No other FPGA board vendor can match the volume of data we can send straight into the heart of the processing elements and then straight back out again.

An FPGA-based high-performance processing engine thrives on data streaming in and out at high rates of speed. The FPGAs should be part of a balanced and unified system architecture, providing maximum performance, with memory, processing power, and I/O speeds designed and integrated for performance, scalability, and growth.

Annapolis Micro Systems, Inc.'s WILDSTAR 4 (Xilinx Virtex-4 based) and WILDSTAR 5 (Xilinx Virtex-5 based) families of FPGA-based processing boards also support an extensive set of extremely high-quality A/D and D/A boards.

Annapolis Micro Systems, Inc. is a world leader in high-performance, COTS FPGA-based boards and processing for RADAR, SONAR, SIGINT, ELINT, DSP, FFTs, communications, Software-Defined Radio, encryption, image processing, prototyping, text processing, and other processing-intensive applications.

Annapolis is famous for the high quality of our products and for our unparalleled dedication to ensuring that the customers' applications succeed.

FEATURES

› Three individually configurable 4X connectors – four lanes per connector

› Up to four 2.5 Gb full duplex Serial FPDP ports per connector

› Up to 25 Gb full duplex RocketIO per connector

› Up to 10 Gb full duplex InfiniBand per connector

› Up to 10 Gb full duplex Ethernet per connector

› Optional onboard oscillators for other line rates like Fibre Channel

› I/O card plugs onto WILDSTAR 4 or 5 VME/VXS/IBM Blade Chassis/PCI-X/PCI Express main board

› JTAG, ChipScope, and Serial Port access

› Proactive thermal management system; available in both commercial and industrial temperature grades

› Includes one-year hardware warranty, software updates, and customer support

› We offer training and exceptional special application development support, as well as more conventional customer support
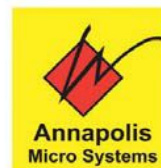
› Full CoreFire Board Support Package for fast, easy application development

› VHDL model, including source code for hardware interfaces

Annapolis Micro Systems, Inc.

190 Admiral Cochrane Drive, Suite 130 • Annapolis, MD 21401 USA
410-841-2514
www.annapmicro.com

Tri XFP I/O Card

Annapolis Micro Systems, Inc. is a world leader in high-performance, COTS FPGA-based boards and processing for RADAR, SONAR, SIGINT, ELINT, DSP, FFTs, communications, Software-Defined Radio, encryption, image processing, prototyping, text processing, and other processing-intensive applications.

The Annapolis Tri XFP I/O Card, which works with the WILDSTAR 4/5 Family Architecture, has three 10 Gb individually configured XFP connectors, each with its own XAUI to XFI converter. Industry-standard pluggable fiber optic transceivers can be purchased from Annapolis or from other vendors. The Tri XFP provides up to 30 Gb full duplex I/O directly between the outside world and the RocketIO pins on the Xilinx Virtex-II Pro or Virtex-4 I/O FPGA on the WILDSTAR 4 main board. No other vendor provides that volume of data straight into the heart of the processing elements and then back out again.

Two I/O cards can reside on each WILDSTAR 4 or WILDSTAR 5 VXS or PCI-X/PCI Express board, with up to 30 million user reprogrammable gates.

The Tri XFP card will support 10 Gb Ethernet, 10 Gb Fibre Channel, and OC-192. Although the protocols will be provided as black box solutions with few modifications by users allowed, more adventurous users who choose to develop their own communications protocols from the basics already have access to all the board resources through VHDL source for the interfaces to SRAM, signal conditioners, LAD bus, I/O bus, and PPC Flash. CoreFire™ users will have the usual CoreFire Board Support Package.

The Tri XFP is the first of many I/O cards Annapolis will be releasing for its new WILDSTAR 4/5 Architecture Family, which uses Xilinx Virtex-4 and Virtex-5 FPGAs for processing elements. WILDSTAR 4 is the tenth generation of Xilinx FPGA processing-based COTS boards from Annapolis.

Annapolis is famous for the high quality of our products and for our unparalleled dedication to ensuring that the customers' applications succeed. We offer training and exceptional special application development support, as well as more conventional customer support.

FEATURES

› Up to 10 Gb Full Duplex Ethernet per connector

› Up to 10 Gb Fibre Channel

› OC-192

› Three 10 Gb XFP connectors

› Accepts industry-standard pluggable transceivers

› Available in both commercial and industrial temperature grades

› Includes one-year hardware warranty, software updates, and customer support

› One or two I/O cards fit on a single WILDSTAR 4/5 processing board

› New I/O form factor for improved thermal performance

› First of many WILDSTAR 4/5 Family I/O cards, including superior performance A/D, D/A, and additional high-speed communication cards

› Save time and effort; reduce risk with COTS boards and software

› Achieve world-class performance; WILD solutions outperform the competition

Annapolis Micro Systems, Inc.

190 Admiral Cochrane Drive, Suite 130 • Annapolis, MD 21401 USA
410-841-2514
www.annapmicro.com
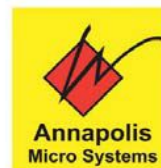
WILDSTAR 5 for IBM Blade

Perfect Blend of Processors and Xilinx Virtex-5 FPGAs. Eleventh Annapolis Generation.

Direct Seamless Connections –
No data reduction between: external sensors and FPGAs, FPGAs and processors over InfiniBand or 10 Gb Ethernet backplane, FPGAs and standard output modules.

Ultimate Modularity –
From zero to six Virtex-5 processing FPGA/memory modules, and two Virtex-5 I/O FPGAs. Accepts one or two standard Annapolis WILDSTAR 4/5 I/O mezzanines: Quad 130 MSps through Quad 500 MSps A/D, 1.5 GSps through 2.2 GSps A/D, Quad 600 MSps DAC, InfiniBand, 10 Gb Ethernet, SFPDP.

Fully Integrated into the IBM Blade Management System –
Abundant power and cooling for maximum performance.

Annapolis Micro Systems, Inc. is a world leader in high-performance, COTS FPGA-based boards and processing for RADAR, SONAR, SIGINT, ELINT, DSP, FFTs, communications, Software-Defined Radio, encryption, image processing, prototyping, text processing, and other processing-intensive applications. We support our board products with a standardized set of drivers, APIs, and VHDL simulation models.

Develop your application very quickly with our CoreFire™ FPGA Application Builder, which transforms the FPGA development process, making it possible for theoreticians to easily build and test their algorithms on the real hardware that will be used in the field. CoreFire, based on dataflow, automatically generates distributed control fabric between cores. Our extensive IP and board support libraries contain more than 1,000 cores, including floating point and the world's fastest FFT. A graphical user interface for design entry supports hardware-in-the-loop debugging and provides proven, reusable, high-performance IP modules.

WILDSTAR 5 for IBM Blade, with its associated I/O cards, provides extremely high overall throughput and processing performance. The combination of our COTS hardware and CoreFire allows our customers to make massive improvements in processing speed, while achieving significant savings in size, weight, power, person-hours, dollars, and calendar time to deployment.

Achieve world-class performance; WILDSTAR solutions outperform the competition.

FEATURES

› From two to eight Virtex-5 FPGA processing elements – LX110T, LX220T, LX330T, FX100T, FX130T, or FX200T; six are pluggable with power module and memory

› Up to 10.7 GB DDR2 DRAM per WILDSTAR 5 for IBM Blade Board

› 144 x 144 crossbar; 3.2 Gb per line; two external PPC 440s – 1 per each I/O FPGA

› Full CoreFire Board Support Package for fast, easy application development

› VHDL model, including source code for hardware interfaces and ChipScope access

› Available in both commercial and industrial temperature grades

› Proactive thermal management system – board-level current measurement and FPGA temperature monitor, accessible through Host API

› Includes one-year hardware warranty, software updates, and customer support

› Blade management controller; USB, RS-485, Ethernet, KVM, 16 RIO, switch to 1 GbE over backplane

› Save time and effort; reduce risk with COTS boards and software

› We offer training and exceptional special application development support, as well as more conventional support

› Famous for the high quality of our products and our unparalleled dedication to ensuring that the customers' applications succeed

---

For more information, contact: wfinfo@annapmicro.com   **www.mil-embedded.com/p35882**

Annapolis Micro Systems, Inc.
190 Admiral Cochrane Drive, Suite 130 • Annapolis, MD 21401 USA
410-841-2514
www.annapmicro.com

WILDSTAR 6 for AMCs

Annapolis Micro Systems, Inc. is a world leader in high-performance, COTS FPGA-based processing for radar, sonar, SIGINT, ELINT, DSP, FFTs, communications, Software-Defined Radio, encryption, image processing, prototyping, text processing, and other processing-intensive applications. Our fourteenth generation WILDSTAR 6 for AMC uses Xilinx's newest Virtex-6 FPGAs for state-of-the-art performance. It accepts one FMC I/O Card. Our boards work on a number of operating systems, including Windows, Linux, Solaris, IRIX, ALTIX, and VxWorks. We support our board products with a standardized set of drivers, APIs, and VHDL simulation models.

Develop your application very quickly with our CoreFire™ FPGA Application Builder, which transforms the FPGA development process, making it possible for theoreticians to easily build and test their algorithms on the real hardware that will be used in the field. CoreFire, based on dataflow, automatically generates distributed control fabric between cores.

Our extensive IP and board support libraries contain more than 1,000 cores, including floating point and the world's fastest FFT. CoreFire uses a graphical user interface for design entry, supports hardware-in-the-loop debugging, and also provides proven, reusable, high-performance IP modules. WILDSTAR 6 for AMC, with its associated I/O Cards, provides extremely high overall throughput and processing performance. The combination of our COTS hardware and CoreFire allows our customers to make massive improvements in processing speed, while achieving significant savings in size, weight, power, person-hours, dollars, and calendar time to deployment.

Annapolis is famous for the high quality of our products and for our unparalleled dedication to ensuring that the customers' applications succeed. We offer training and exceptional special application development support, as well as more conventional support.

Save time and effort and reduce risk with our COTS boards and software. Achieve world-class performance – WILD solutions outperform the competition.

FEATURES

› One Xilinx Virtex-6 FPGA I/O Processing Elements – LX240T, LX365T, LX550T, SX315T or SX475T

› On board Host Freescale P1020 or P2020 PowerPC

› Up to 2.5 GBytes DDR2 DRAM in 5 memory banks or

› Up to 80 MB DDRII or QDRII DRAM in 5 memory banks

› Programmable FLASH to store FPGA image

› 4X PCI Express Bus Gen 2 between PPC and FPGA

› Supports VITA 57 FMC I/O Cards

› Full CoreFire Board Support Package for fast, easy application development

› VHDL model, including source code for hardware interfaces and ChipScope access

› Available in both commercial and industrial temperature grades

› Proactive Thermal Management System – current, voltage, and temperature monitoring sensors via Host API

› Includes one year hardware warranty, software updates, and customer support. Training available.

For more information, contact: wfinfo@annapmicro.com

Annapolis Micro Systems, Inc.

190 Admiral Cochrane Drive, Suite 130 • Annapolis, MD 21401 USA
410-841-2514
www.annapmicro.com

WILDSTAR 6 PCIe

Annapolis Micro Systems, Inc. is a world leader in high-performance, COTS FPGA-based processing for radar, sonar, SIGINT, ELINT, DSP, FFTs, communications, Software-Defined Radio, encryption, image processing, prototyping, text processing, and other processing-intensive applications. Our fifteenth generation WILDSTAR 6 for PCI Express uses Xilinx's newest Virtex-6 FPGAs for state-of-the-art performance. It accepts one or two I/O mezzanine cards, including Single 1.5 GHz 8 Bit ADC, Quad 250 MHz 12 Bit ADC, Single 2.5 GHz 8 Bit ADC, Quad 130 MHz 16 Bit ADC, Dual 2.3/1.5 GSps 12 Bit DAC, Quad 600 MSps 16 Bit DAC, Universal 3Gbit Serial I/O (RocketIO, 10 Gb Ethernet, InfiniBand), and Tri XFP (OS 192, 10G Fibre Channel, 10 Gb Ethernet). Our boards work on a number of operating systems, including Windows, Linux, Solaris, IRIX, ALTIX, and VxWorks. We support our board products with a standardized set of drivers, APIs, and VHDL simulation models.

Develop your application very quickly with our CoreFire™ FPGA Application Builder, which transforms the FPGA development process, making it possible for theoreticians to easily build and test their algorithms on the real hardware that will be used in the field. CoreFire, based on dataflow, automatically generates distributed control fabric between cores.

Our extensive IP and board support libraries contain more than 1000 cores, including floating point and the world's fastest FFT. CoreFire uses a graphical user interface for design entry, supports hardware-in-the-loop debugging, and also provides proven, reusable, high-performance IP modules. WILDSTAR 6 for PCI Express, with its associated I/O Cards, provides extremely high overall throughput and processing performance. The combination of our COTS hardware and CoreFire allows our customers to make massive improvements in processing speed, while achieving significant savings in size, weight, power, person-hours, dollars, and calendar time to deployment.

Annapolis is famous for the high quality of our products and for our unparalleled dedication to ensuring that the customers' applications succeed. We offer training and exceptional special application development support, as well as more conventional support.

Save time and effort and reduce risk with COTS boards and software. Achieve world-class performance – WILD solutions outperform the competition.

FEATURES

> Up to three Xilinx Virtex-6 FPGA I/O processing elements – LX240T, LX365T, LX550T, SX315T, or SX475T

> Up to 8 GBytes DDR2 DRAM or DDR3 DRAM in 14 memory banks per WILDSTAR 6 for PCI Express board or up to 480 MBytes DDRII+/QDRII DRAM in 15 memory banks

> Programmable FLASH for each FPGA to store FPGA images

> 8X PCI Express Bus Gen 1 or Gen 2

> Supports PCI Express standard external power connector

> High-speed DMA Multi-Channel PCI controller

> Full CoreFire Board Support Package for fast, easy application development

> VHDL model, including source code for hardware interfaces and ChipScope access

> Available in both commercial and industrial temperature grades

> Proactive Thermal Management System – Board Level current measurement and FPGA temperature monitor, accessible through Host API

> Includes one year hardware warranty, software updates, and customer support. Training available.

Annapolis Micro Systems, Inc.

190 Admiral Cochrane Drive, Suite 130 • Annapolis, MD 21401 USA
410-841-2514
www.annapmicro.com

WILDSTAR 6 – Pluggable Virtex-6 Module for IBM Blade

Annapolis Micro Systems is a world leader in high-performance, COTS FPGA-based boards and processing for RADAR, SONAR, SIGINT, ELINT, DSP, FFTs, communications, Software-Defined Radio, encryption, image processing, prototyping, text processing, and other processing-intensive applications.

Our 14th-generation Pluggable Virtex-6 Module for WILDSTAR 5 for IBM BladeCenter uses Xilinx's newest Virtex-6 FPGAs for state-of-the-art performance. This module plugs into the WILDSTAR 5 for IBM BladeCenter. Other potential modules available today include Virtex-5 and Tilera. We support our board products with a standardized set of drivers, APIs, and VHDL simulation models.

Develop your application very quickly with our CoreFire™ FPGA Application Builder, which transforms the FPGA development process, making it possible for theoreticians to easily build and test their algorithms on the real hardware that will be used in the field. CoreFire, based on dataflow, automatically generates distributed control fabric between cores.

Our extensive IP and board support libraries contain more than 1,000 cores, including floating point and the world's fastest FFT. CoreFire uses a graphical user interface for design entry, supports hardware-in-the-loop debugging, and provides proven, reusable, high-performance IP modules. The Virtex-6 Pluggable Module provides extremely high overall throughput and processing performance. The combination of our COTS hardware and CoreFire allows our customers to make massive improvements in processing speed, while achieving significant savings in size, weight, power, person-hours, dollars, and calendar time to deployment.

Annapolis is famous for the high quality of our products and for our unparalleled dedication to ensuring that the customers' applications succeed. We offer training and exceptional special application development support, as well as more conventional support.

FEATURES

› One Virtex-6 FPGA processing element – XC6LX240T, XC6LX365T, XC6LX550T, XC6SX315, or XC6SX475
› Up to 3.5 GB DDR2 DRAM in 7 banks or up to 224 MB DDRII or QDRII SRAM
› Plugs into any of 6 pluggable processing module slots on WILDSTAR 5 for IBM BladeCenter
› Up to 16 RocketIO lanes to crossbar
› 4 lanes connect to PCIe switch on WILDSTAR 5 Blade Board
› LVDS systolic ring connecting all I/O FPGAs and computing FPGAs
› RocketIO systolic ring connecting pluggable positions 0-2 and 3-5
› Programmable Flash to store FPGA images on WILDSTAR 5 Blade Board, which acts as host
› Full CoreFire Board Support Package for fast, easy application development
› VHDL model, including source code for hardware interfaces and ChipScope access
› Host software: Windows, Linux, VxWorks, etc.
› Available in both commercial and industrial temperature grades
› Proactive thermal management system – board-level current measurement and FPGA temperature monitor, accessible through Host API
› Save time and effort and reduce risk with COTS boards and software
› Achieve world-class performance – WILD solutions outperform the competition
› Includes one-year hardware warranty, software updates, and customer support; training available

---

For more information, contact: wfinfo@annapmicro.com

Annapolis Micro Systems, Inc.

190 Admiral Cochrane Drive, Suite 130 • Annapolis, MD 21401 USA
410-841-2514
www.annapmicro.com

WS4 Quad 250/400/500 MSps A/D

The Annapolis Quad Channel 250/400/500 MSps A/D I/O Card provides 4 A/D inputs with converter speeds of up to 250, 400, or 500 MHz and resolutions of 13, 14, or 12 bits, respectively. The board has four A/D converters from TI (ADS5444, ADS5474, or ADS5463) fed by onboard analog input circuits that convert the single-ended, 50-ohm SMA input into differential signals for the ADC.

There is an onboard ultra-low jitter and skew clock distribution circuit to allow all four channels on a single A/D I/O board to be synchronized together. There is also an external clock input and a trigger input allowing multiple A/D I/O cards to be synchronized together. Synchronization of A/D I/O cards can be facilitated by the Annapolis 4 or 8 Channel Clock Distribution Boards.

In concert with the WILDSTAR 4 or WILDSTAR 5 FPGA processing main boards, this mezzanine board supplies user-configurable real-time continuous sustained processing of the full data stream. Up to two A/D I/O cards can reside on each WILDSTAR 4 or WILDSTAR 5 VME/VXS or IBM Blade main board or reside on one A/D I/O card on each PCI-X or PCIe main board.

Annapolis Micro Systems, Inc. is a world leader in high-performance, COTS FPGA-based boards and processing for RADAR, SONAR, SIGINT, ELINT, DSP, FFTs, communications, Software-Defined Radio, encryption, image processing, prototyping, text processing, and other processing-intensive applications.

Our boards run on many different operating systems. We support our boards with a standardized set of drivers, APIs, and VHDL simulation models. VHDL source is provided for the interfaces to A/Ds, D/As, DRAM/SRAM, LAD bus, I/O bus, and PPC Flash. CoreFire™ users will have the usual CoreFire Board Support Package.

The combination of our COTS hardware and our CoreFire FPGA Application Development tool allows our customers to make massive improvements in processing speed, while achieving significant savings in size, weight, power, person-hours, dollars, and calendar time to deployment.
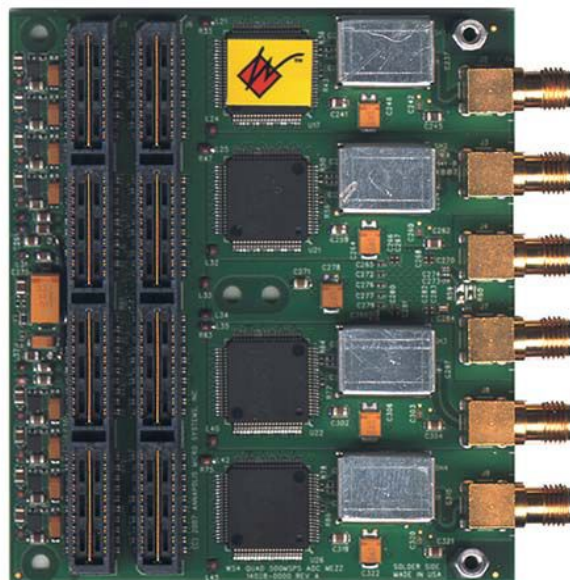
FEATURES

› Four TI A/D converters of one of the speed and bit size types: ADS5444 250 MSps 13 bits, ADS5474 400 MSps 14 bits, ADS5463 500 MSps 12 bits

› Analog input bandwidths of up to: 500 MHz for the 250 MSps A/D board, 1,400 MHz for the 400 MSps A/D board, 2,000 MHz for the 500 MSps A/D

› Six SMA front panel connectors: four 50-ohm analog inputs, one single-ended 50-ohm clock input, one trigger input

› Onboard ultra-low jitter and skew clock distribution circuit to allow synchronization of all four channels on a single I/O card

› I/O card plugs onto WILDSTAR 4 or 5 VME/VXS/PCI-X/PCIe/ IBM Blade main boards

› JTAG, ChipScope, and Serial Port access

› Proactive thermal management system; available in both commercial and industrial temperature ranges

› Full CoreFire Board Support Package for fast, easy application development and technology refresh

› VHDL model, including source code for hardware interfaces

› Includes one-year hardware warranty, software updates, and customer support; reduce risk with COTS

› We offer training and exceptional special application development support, as well as more conventional customer support

› Annapolis is famous for the high quality of our products and for our unparalleled dedication to ensuring that the customers' applications succeed

---

For more information, contact: wfinfo@annapmicro.com

Acromag

30765 S. Wixom Road • Wixom, MI 48393 USA
248-295-0310
www.acromag.com

VPX cards with configurable Spartan-6 or Virtex-5 FPGA

Acromag's new 3U VPX FPGA boards provide powerful and cost-effective solutions for high-speed processing of complex algorithms in embedded computing applications. The VPX-VLX model offers a choice of logic-optimized Xilinx Virtex-5 FPGAs while the VPX-SLX version employs a Spartan-6 FPGA with 150k logic cells. Virtex-5 FPGA cards deliver maximum computing performance. Spartan-6 FPGA cards are designed for cost-sensitive mid-level applications. Both cards feature a high-throughput PCI Express interface, generous dual-ported memory for efficient data handling, and 64 I/O lines direct to the FPGA. Ideal for defense, aerospace, or scientific research; typical applications involve signal intelligence, image processing, and hardware simulation.

A double fat pipe 4-lane PCIe interface ensures very fast data throughput. Dual-port SRAM enables high-speed DMA transfers to/from the CPU while simultaneously writing data to memory. Large, DDR2 SDRAM buffers give the FPGA fast access to I/O port data.

FEATURES

> Virtex-5 or Spartan-6 configurable FPGA with up to 150k logic cells
> PCI Express 4-lane bus interface
> Supports both front and rear I/O (64 I/O or 32 LVDS via rear J4/P4)
> Plug-in I/O interface cards for front mezzanine (14-bit A/D; differential, TTL, LVDS digital I/O)
> Two banks 1M x 32-bit dual-port SRAM
> Two banks 32M x 16-bit DDR2 SDRAM
> Dual DMA channel data transfers
> FPGA code loads from the PCIe bus or from on-board flash memory
> Air-cooled (0 to 70°C) and conduction-cooled (-40 to 85°C) models

For more information, contact: solutions@acromag.com   www.mil-embedded.com/p52925

CPU Tech

5731 West Las Positas Blvd. • Pleasanton, CA 94588 USA
925-224-9920
www.cputech.com

Acalis CPU872-G02B Secure Processor

CPU Tech's CPU872 Secure Processor contains dual high-performance PowerPC® 440 and FPU cores, dual embedded DRAMs and numerous other functions. It provides maximum sustained performance on real-world problems while consuming only 8 watts. The CPU872 also contains intrinsic features that prevent tampering and reverse-engineering to protect valuable intellectual property contained in the software and system. The target markets include financial, military and aerospace, communications and others that maintain critical information systems at peak sustained performance.

To guard against the insertion of malicious circuits or other Trojan Horses, the CPU872 is fabricated at the IBM Trusted Foundry in cooperation with the Trusted Access Program Office. The secure processing architecture contains a multi-layered group of effective and resilient run-time techniques so the user can select the desired level of protection. In all cases, the protection does not impact the performance of the device. For additional security details, contact CPU Tech.
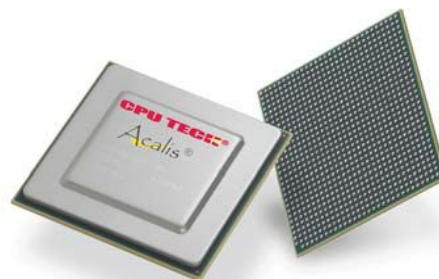
FEATURES

> In addition to the dual 2.0 DMIPS/MHz PowerPC and dual 64-bit, 2.0 MFLOPS/MHz FPU cores, the CPU872 also includes utility compute engines that accelerate I/O performance off-loading the 440 cores. These utility engines ensure that the 440 and FPU cores operate at peak sustained performance so applications complete in less time and use minimal power compared to other PowerPC-based offerings.

> The Acalis Secure Processor and documentation are available today, with full security configuration tool suite available to designers in early 2012.

For more information, contact: acalis@cputech.com   www.mil-embedded.com/p45793

Innovative Integration

2390-A Ward Avenue • Simi Valley, CA 93065 USA
www.innovative-dsp.com

X6-1000M

The X6-1000M integrates high-speed digitizing and signal generation with signal processing on a PMC/XMC I/O module for demanding DSP applications. The tight coupling of the digitizing to the Virtex-6 FPGA core realizes architectures for SDR, RADAR, and LIDAR front-end sensor digitizing and processing. The PCI Express system interface sustains transfer rates over 2 GB/s for data recording and integration as part of a high-performance, real-time system. Available FPGAs: Xilinx Virtex-6 SX315T/SX475T or LX240T.

The X6-1000M features two, 12-bit 1 GSPS A/Ds and two 1 GSPS 16-bit DACs. Analog input bandwidth of over 2 GHz supports wideband applications and undersampling. The DACs have features for interpolation and coarse mixing for upconversion. The sample clock is from either a low-jitter PLL or external input. Multiple cards can be synchronized for sampling and down-conversion.

FEATURES

> Two 1 GSPS, 12-bit A/D channels
> Two 1 GSPS, 16-bit DAC channels
> ±1V, AC or DC-Coupled, 50 ohm, SMA inputs and outputs
> 4 Banks of 1 GB DRAM (4 GB total)
> Ultra-low-jitter programmable clock
> Arbitrary Waveform Generation Memory
> Gen2 x8 PCI Express providing 2 GB/s sustained transfer rates
> PCI 32-bit, 66 MHz with P4-to-Host card
> 20-25W typical Conduction Cooling per VITA 20
> Ruggedization Levels for Wide Temperature Operation
> Adapters for VPX, CompactPCI, desktop PCI, and cabled PCI Express systems

For more information, contact: sales@innovative-dsp.com   **www.mil-embedded.com/p52897**

Xilinx

2100 Logic Drive • San Jose, CA 95124 USA
408-559-7778
www.xilinx.com

Defense-grade Spartan®-6Q FPGA family

The Defense-grade Spartan®-6Q FPGA family is ideal for MILCOM and other applications where security, low power and low costs are paramount.

Advanced Power Management and superior integration provide the optimal balance of size, weight, power and cost (SWaP-C), while 3rd generation Information Assurance and Anti-tamper capabilities – including Security Monitor (SECMON) IP and Type-1 Single-Chip Cryptography – help designers comply with regulations such as the U.S. Department of Defense (DoD)-mandated 5000 Series AT requirements.

FEATURES

> 3rd generation Information Assurance (Single-Chip Cryptography Type-1)
> 3rd generation Anti-tamper (DoD 5000 Series)
> Superior SWaP-C with small form factor (SFF) packaging, Single-Chip Cryptography and integrated external components
> Lowest overall cost defense-grade FPGA with DDR3 support
> True, full-range tested I and Q temperatures
> Pin-compatible commercial-grade portfolio
> Long term product support
> Mask set control

For more information, contact: moreinfo@xilinx.com   **www.mil-embedded.com/p53156**

Xilinx

2100 Logic Drive • San Jose, CA 95124 USA

408-559-7778

www.xilinx.com

Defense-grade Virtex®-6Q FPGA family

The Defense-grade Virtex®-6Q FPGA family offers advancements in secure performance and reliability – built with the right mix of programmability, integrated blocks for DSP memory and connectivity support – to satisfy the insatiable demand for higher bandwidth and higher performance with higher security. Third generation Anti-tamper features, including Security Monitor (SECMON) IP, uniquely give designers the latest security advancements while simultaneously supporting DO-254 enablement, making Virtex-6Q devices the right fit for a broad range of applications.

The Virtex-6Q FPGA family also features ruggedized packaging protection against "tin-whiskering" and caustic solvent cleaning systems, and is fully tested for extended temperature ranges, giving designers a choice of solutions qualified for operation in either I-temperature or M-temperature ranges.

FEATURES

> 3rd generation Anti-tamper (DoD 5000 Series)
> Superior SWaP-C and integration
> True, full range-tested I and M temperatures
> Ruggedized packaging
> Off-the-shelf availability
> Pin-compatible commercial-grade portfolio
> 16-year, long-term product support
> Mask set control

For more information, contact: moreinfo@xilinx.com   **www.mil-embedded.com/p53157**

Xilinx

2100 Logic Drive • San Jose, CA 95124 USA

408-559-7778

www.xilinx.com

Space-grade Virtex®-5QV FPGA

The Space-grade Virtex®-5QV FPGA is the industry's first high-performance rad-hard reconfigurable FPGA for processing-intensive space systems. Available off-the-shelf, the device offers the highest density, performance and integration capabilities, enabling more complex and capable systems over rad-hard ASIC devices with their high development costs and long lead times, or traditional one-time programmable (OTP) solutions.

FEATURES

> SEU Latch-up Immunity (LET$_{TH}$)
  – Configuration Cell Upset Rate (GEO)
  – Functional Interrupt Rate (GEO)
  – Total Ionizing Dose
  – Dose Rate Upset
  – Dose Rate Latch-up
  – > 100 Mev-cm^2/mg
  – < 3.8 x 10^{-10} Upsets/Bit-Day
  – < 10^{-10} Upsets/Bit-Day
  – > 1 Mrad(Si)
  – > 10^9 Rad(Si)/s
  – > 10^{10} Rad(Si)/s

> Up to 450MHz ultra-high performance DSP technology with flexible embedded processing
> Up to 130,000 logic cells
> 836 user I/Os programmable for easy integration
> Integrated SERDES solution with 18 channels of 3GHz multi-gigabit serial transceivers
> SelectIO™ technology
> 10.7Mb of powerful 36Kbit block RAM/FIFOs
> 2nd-generation 25x18 DSP slices
> PCI Express™-compliant integrated Endpoint blocks

> No NREs for mask sets
> No design/manufacturing cycle time
> Pin compatible to commercial Virtex-5 FX130T FPGA with commercial prototyping board available
> Daisy chain part available
> Xilinx V-grade flow (ADQ-0007)
> Ruggedized ceramic column grid array packaging

For more information, contact: moreinfo@xilinx.com   **www.mil-embedded.com/p53158**

WinSystems, Inc.

715 Stadium Drive • Arlington, TX 76011 USA
817-274-7553
www.WinSystems.com

Fanless 1.66GHz Intel® Atom™ EPIC SBC

WinSystems' EPX-C380 is a full-featured EPIC-compatible SBC with a rich array of onboard peripherals such as video, Ethernet, USB, and serial, plus more I/O expansion options with PC/104 and MiniPCIe modules.

It also supports either the Intel® Atom™ single core N450 or dual core D510 processors. Since the EPX-C380's architecture is PC-compatible, it supports Windows® XP Embedded, WES7, and Linux software operating systems, along with a vast software development tool set including device drivers and libraries.

The board is designed for rugged applications, including industrial automation, security, Mil/COTS, and transportation. It operates over a temperature range of -40°C to +70°C without a fan or the necessity to slow down the CPU clock frequency.

For more information go to
www.winsystems.com/EPX-C380MR

FEATURES

> Intel® Atom™ 1.66GHz single- or dual-core CPU
> Advanced power management features
> Supports CRT and flat panels simultaneously
> Two Intel Gigabit Ethernet controllers
> 802.11a/b/g wireless supported with MiniPCIe
> Four serial COM ports, four USB 2.0 ports, and 48 bidirectional TTL digital I/O lines
> Two SATA channels and CompactFlash supported
> PC/104 and MiniPCIe module expansion
> Starter kits to speed system development

For more information, contact: Info@WinSystems.com

www.mil-embedded.com/p47868

COTS Collection: Boards, Carriers, and Mezzanines: General Purpose I/O

ACCES I/O Products, Inc.

10623 Roselle Street • San Diego, CA 92121 USA
858-550-9559
www.accesio.com

USB-DIO-96 High-Density Digital I/O

Designed for compact control and monitoring applications, this product features 96 or 48 industrial-strength TTL digital I/Os. This USB device is an ideal solution for adding portable, easy-to-install digital I/O to any PC or embedded system with a USB port. The USB-DIO-96 is useful for monitoring dry contacts or generating outputs for controlling external devices such as LEDs and other indicators or system equipment. Applications include home, portable, laptop, education, laboratory, industrial automation, and embedded OEM.

These boards use 2 or 4 industry-standard 50-pin IDC-type shrouded headers with 24 lines per connector. Utility 5VDC is available on pin 49 of each connector, with grounds on all even-numbered pins to reduce crosstalk and maintain industry compatibility. A mini USB header connector is provided in parallel with the high-retention Type B connector for stacking and embedded applications. Available accessories include a wide variety of cables and screw terminal boards for quick and easy connectivity.

FEATURES

> 96 or 48 lines of digital I/O
> High-speed USB 2.0 device, USB 1.1 backward-compatible
> Twelve or six 8-bit ports independently selectable for inputs or outputs
> All I/O lines buffered with 32mA source, 64mA sink current capabilities
> I/O buffers can be enabled or tri-stated under program control
> Jumper-selectable I/O pulled up to 5V for contact monitoring, down-to-ground, or floating
> Resettable fused +5VDC outputs per 50-pin connector
> OEM version (board only) features PC/104 module size and mounting compatibility

For more information, contact: contactus@accesio.com

www.mil-embedded.com/p42373

ALPHI Technology Corporation

1898 E. Southern Ave. • Tempe, Arizona 85282 USA
480-838-2428
www.AlphiTech.com

XPC-1553-S

The latest member of the ALPHI 1553 product line is a stand-alone, fully sealed, rugged, ultra-small and ultra-low-power Intel Atom™-based module. The system is a fully submersible unit with an ultra-small-footprint that is only 5.25" x 3.5" and includes an SBC and mass storage.

Dual-redundant MIL-STD-1553 Interface, programmable Bus Controller, Remote Terminal, or Bus Monitor modes, supports MIL-STD-1553A/B. Controlled by a 1.6GHz Atom™ processor, 512KB of L2 Cache and Up to 2GB of 533MHz DDR-2 SDRAM. The system I/O configuration is based on two non-switched x1 PCIExpress lanes used on two Express Mini Cards slots.

The 1553 capabilities can be complimented by ARINC-429, FireWire, CANbus, GPS, GigE, A/D, D/A or any other functionality.
• Support for Windows® XP/XPE/7, Linux® and VxWorks®
• Available in extended temp -40°C to +85°C

This is a perfect solution for a wide array of 1553 communication applications such as: military mission computer, ground vehicle and many others.

FEATURES

> Ultra-small only 5.25" x 3.5" SBC that is a Fully Sealed Submersible unit with MIL-STD-810F Connector
  – 1.6GHz Atom™ processor with 512KB of L2 Cache
  – Up to 2GB of 533MHz DDR-2 SDRAM
  – High-performance graphics with 3D acceleration
  – 24-bit LVDS for direct connection to LCD displays
> Programmable 1553 mode: BC, RT and BM
  – Eight buffered General Purpose I/O lines
  – Support for one SSD
  – Support for Express Mini Cards
  – One Fast Ethernet port
  – Five USB 2.0 ports, one Serial port with RS232/422 support
  – Onboard Power Supplies for single 5-12VDC input

For more information, contact: sales@Alphitech.com

www.mil-embedded.com/p53151

Kontron

14118 Stowe Drive • Poway, CA 92064-7147 USA
888-294-4558
www.kontron.com

Kontron COM Express Modules

Kontron offers the most expansive portfolio of COM Express-compliant and compatible modules. Modules range in performance and power consumption from the ultra-small, extremely power efficient nanoETXexpress-TT with the newest Intel Atom Exxx processor series all the way up to the top-performing ETXexpress-SC featuring second-generation Intel Core i7 processors.

For more information on all of the Kontron COM Express solutions, visit www.kontron.com/ETXexpress.

FEATURES

> COM Express Type 1, Type 2, Type 6 and Type 10 solutions
> Processing performance up to 2.53GHz Intel® Core™ i7 processors
> Integrated advanced graphics capability with the AMD G-Series APUs and the Fusion Controller Hub
> Other low-power, CPU-based modules also available
> Up to 16MB system memory on COMs with dual-channel and ECC memory support
> ETXexpress, microETXexpress, nanoETXexpress – size options all following the COM Express standard pin-outs
> Request a sample today and start evaluating immediately

For more information, contact: info@us.kontron.com

www.mil-embedded.com/p42374

Advanced Digital Logic

4411 Morena Blvd., Suite 101 • San Diego, CA 92117 USA
858.490.0597
www.adl-usa.com

ADLQM67PC PCIe/104 Intel Core i7/i5 – 2.2 GHz - 2.5 GHz SBC

The ADLQM67PC features the 2nd generation Intel Core™ i7 and i5 processors that incorporate Intel's latest embedded two-chip platform. These second-generation i5 and i7 processors integrate Intel's HD Graphics 3000 engine with AVX (advanced vector extensions) as well as the memory controller functions of a traditional GMCH.

The QM67 Platform Controller Hub (PCH) provides PCI Express I/O bandwidth at twice the speed (5GT/s) of previous i5/i7 or Core 2 Duo platforms to bring unparalleled performance to applications such as radar and sonar processing, image signal processing, surveillance and reconnaissance and UAVs. This is further enhanced through the ability to do intensive general-purpose computation on the Graphics Processing Units (GPGPU). Serving the PC/104 Industry for more than 15 years!

Shipping Q4!

FEATURES

> Up to 4GB DDR3-1333, DRAM, SODIMM204
> Type 1 Bottom-Stacking PCI/104 Express V2.0 with Gen2 protocol (2.5 to 5GT/s)
> 2x SATA 600 Ports with RAID
> 2x 10/100/1000Mbit Ethernet LAN Port
> 2x RS-232 COM Ports, 8x USB2.0 Ports
> Separate Onboard VGA, LVDS, HDMI / DVI, DisplayPort
> ADL-SST (System Sensor Technology)
> Windows XP, XPe, Windows 7, Compact Embedded Support
> Linux Compatible
> Operating Range: -25° to +70º C (-40º to +85º C available – call sales for details)

For more information, contact: sales@adl-usa.com   **www.mil-embedded.com/p53141**

WinSystems, Inc.

715 Stadium Drive • Arlington, TX 76011 USA
817-274-7553
www.WinSystems.com

-40°C to +85°C PC/104 SBC with Dual Ethernet

The PCM-VDX-2 is a highly integrated PC/104-compatible Single Board Computer (SBC) designed for space-limited and low-power applications.

The full-featured SBC includes a 1GHz Vortex86DX with 512KB of DRAM. I/O support includes two 10/100 Ethernet ports, four USB 2.0 ports and four COM channels. Its low power dissipation permits fanless operation over a temperature range from -40°C to +85°C. This board is well suited for rugged applications requiring excellent processor performance in an embedded PC design.

The PCM-VDX-2 has x86 PC software compatibility, which assures a wide range of tools to aid in your application's program development and checkout. It supports both Windows® XP Embedded and Linux operating systems and other real-time operating systems. WinSystems provides free technical phone support to assist customers with system integration of our SBCs and I/O modules with their designs.

FEATURES

> Fanless 1GHz Vortex86DX processor
> Small size: 90mm x 96mm
> Two 10/100 Mbps Ethernet controllers
> Four USB 2.0 ports with overcurrent protection
> Four serial RS-232/422/485 channels with FIFOs
> 16 digital I/O lines with event sense supported
> AC'97 audio, PATA, LPT, mouse and keyboard controllers
> MiniPCI and PC/104 expansion connectors
> -40°C to +85°C operating temperature
> Long-term PC/104 product availability

For more information, contact: Info@WinSystems.com   **www.mil-embedded.com/p52377**

Military Embedded Systems   Resource Guide

VersaLogic Corp.
4211 W. 11th Ave. • Eugene, OR 97402 USA
800-824-3163
www.VersaLogic.com/Tiger

VERSALOGIC
C O R P O R A T I O N

Tiger

The Tiger is an embedded single board computer (SBC) featuring a high-performance Intel Atom Z5xx processor. Based on the PC/104-*Plus* industry standard form factor, the Tiger supports both PC/104™ and PC/104-*Plus* stackable expansion boards. With its combination of high performance (up to 1.6 GHz), low power consumption (6W typ. while executing code), and fanless operation, the Tiger is an ideal embedded computer solution for size, weight, and power (SWaP)-constrained embedded applications in industrial, energy, defense/aerospace, medical, and robotics markets.

Driven by an Intel Atom Z5xx processor designed specifically for embedded applications, the Tiger runs completely fanless at up to 1.6 GHz (commercial temperature) or 1.33 GHz (industrial temperature). Enhanced Intel SpeedStep® Technology provides dynamic processor frequency scaling to meet instantaneous performance needs while minimizing power draw and heat dissipation. This allows users to fine-tune the balance between power conservation and performance to suit their application needs. Enhanced low-power states, including the new C6 state (Deep Power Down Technology), allow designers to further minimize overall power consumption.

The integrated graphics core of the Atom Z5xx processor supports advanced 3D graphics and high-definition video decode. Video output is provided through an integrated LVDS flat panel video interface and optional analog VGA support.

Tiger's standard on-board features include Gigabit Ethernet with network boot capability, SO-DIMM socket for up to 2 GB DDR2 RAM, seven USB 2.0 ports, four serial ports, IDE controller with support for two devices, HD audio, and a Disk-on-Module (DOM) socket for removable flash storage. PC/104-*Plus* expansion provides plug-in access to a wide variety of industry-standard expansion modules. VersaLogic's SPX expansion interface provides access to cost-effective plug-in analog, digital, CANbus, and custom I/O solutions.

Available in both industrial (-40° to +85°C) and commercial (0° to +60°C) temperature versions, the Tiger meets MIL-STD-202G specifications for shock and vibration. Transient voltage suppression (TVS) devices on critical I/O ports provide enhanced electrostatic discharge (ESD) protection, which is critical in many OEM applications.

The Tiger features an embedded BIOS with OEM enhancements from Phoenix Technologies. The field-reprogrammable BIOS supports custom defaults and the addition of firmbase applications for security processes, remote booting, and other pre-OS software functions. The Tiger is compatible with Windows, Windows Embedded, Linux, VxWorks, and QNX.
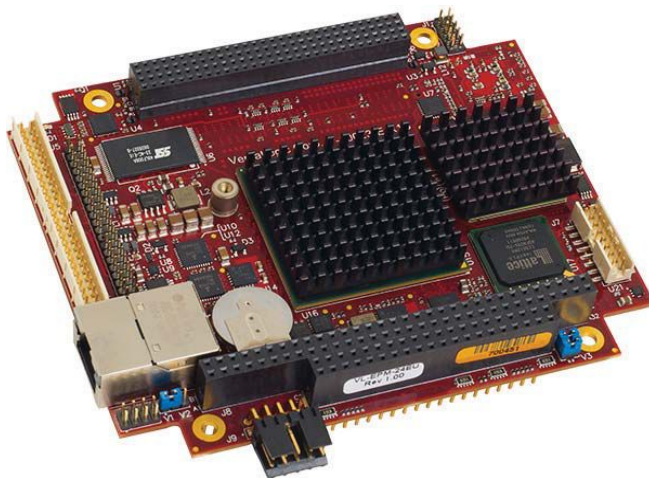
FEATURES

> **PC/104-*Plus* Form Factor**
  Supports PC/104 and PC/104-*Plus* expansion modules in a highly rugged format.

> **Intel Atom Z5xx Processor**
  Up to 1.6 GHz performance with very low power draw.

> **High-performance Video**
  Advanced 3D graphics and high-definition video decode.

> **Network Support**
  Gigabit Ethernet with remote boot support.

> **System RAM**
  Up to 2 GB DDR2 RAM.

> **USB I/O**
  Seven USB 2.0 ports support keyboard, mouse, and other devices.

> **Device I/O**
  Four serial ports, IDE interface, and HD audio.

> **Flash Memory**
  Disk-on-Module socket for plug-in flash storage.

> **Fanless Operation**
  No moving parts required for CPU cooling.

> **Industrial Temperature Version**
  -40° to +85°C operation for harsh environments.

> **MIL-STD-202G**
  Qualified for high shock/vibration environments.

> **SPX Expansion**
  Add additional analog, digital, or CANbus modules.

For more information, contact: Sales@VersaLogic.com   **www.mil-embedded.com/p52216**

WinSystems, Inc.

715 Stadium Drive • Arlington, TX 76011 USA
817-274-7553
www.WinSystems.com

Multifunction PC/104 A/D, D/A, and DIO Module

The PCM-MIO is a versatile, PC/104-based analog input, analog output, and digital I/O board. The board's precision converters and voltage references require no calibration.

It will support up to 16 single-ended input channels or 8 differential input channels with a 16-bit A/D. The software-programmable input ranges are ±5V, ±10V, 0-5V, and 0-10V. There are eight, 12-bit Digital-to-Analog (D/A) converters with individual software programmable voltage ranges of ±5V, ±10V, 0-5V, and 0-10V. A total of 48 bidirectional TTL-compatible digital I/O lines can be software configured as input, output, or output with readback.

The PCM-MIO operates over the industrial temperature range of -40°C to +85°C. Free software drivers are available for C, Windows®, and Linux.

WinSystems also offers one configuration with A/D only and another with only D/A to reduce system cost.

FEATURES

> Two, 16-bit 100K samples/sec Analog-to-Digital (A/D) converters
> Two quad, 12-bit Digital-to-Analog (D/A) converters
> 48 bidirectional TTL-compatible digital I/O lines
> Free software drivers in C, Windows®, and Linux
> -40°C to +85°C operational temperature
> Lower cost A/D- or D/A-only configurations available
> 30-day product evaluation program
> Long term availability

For more information, contact: Info@WinSystems.com

www.mil-embedded.com/p41298

Pinnacle Data Systems, Incorporated

6600 Port Road • Groveport, OH 43125 USA
Tel: +1 (614) 748-1150 • Fax: +1 (614) 748-1209
www.pinnacle.com

PMC-SD18 and XMC-SD18 SATA HDD/SSD Storage Modules

These new SATA Storage Modules are offered in both PMC and XMC formats. Both provide high-capacity SATA storage using compact 1.8 inch hard disk drives (HDD) or solid state drives (SSD) – up to 160GB of storage is available with either drive type. Whether configured with an economical rotating HDD or with a highly shock-resistant SSD, these low-profile modules fit comfortably into VITA 42.3-compatible VME, CompactPCI®, AdvancedTCA®, and PCI Express processor boards without risk of mechanical interference.

The onboard 4-port SATA controller provides 3 additional external drive interfaces. OS support includes Windows, Linux, Solaris x86, and Solaris SPARC. Critical military and aerospace applications will appreciate the high operating shock resistance (1000+ G) and high MTBF (over 1 million hours) when configuring these modules with the latest SSD technology.

FEATURES

> High-capacity 1.8 inch SATA storage PMC and XMC
> Low-cost rotating HDDs for normal use
> Rugged SSDs available for high shock and vibration
> Up to 160GB SSD storage capacity
> Featuring Intel advanced SSD technology (80GB and up)
> 3 additional SATA channels
> Windows, Linux, and Solaris support
> Customization welcomed; extended availability assured

For more information, contact: info.sales@pinnacle.com

www.mil-embedded.com/p41871

Aitech Defense Systems

19756 Prairie Street • Chatsworth, CA 91311 USA
888-248-3248
www.rugged.com/defensepmcs.htm

M59x Series Graphics Boards

Aitech has expanded its family of M59x graphics boards designed for a wide variety of rugged avionics applications with the M595 PMC and M597 XMC.

These high-performance, rugged boards simultaneously drive two independent video streams in a wide variety of graphics and output formats for flexible video input and frame-grabbing formats to meet users' specific application needs.

Both single-width mezzanine boards integrate multiple supporting 2D/3D hardware engines, such as graphics language accessories, parallel processing engines and video and audio decompression units, and complex 3D shading algorithms to support the most advanced graphics and video format standards. This includes LVDS, SDI, HDI, SMPTE 292 and H.264, and graphics languages including DirectX, OpenGL and OpenCL.

The new M595 and M597 use the advanced AMD/ATI E4690 graphics processing unit (GPU) operating at 600 MHz with a 512 MB on-chip GDDR3 SDRAM frame buffer. The E4690, which enables multiple video outputs from its native video ports eliminating the need for external transmitters or encoders, works with an integrated, on-board FPGA to support a wide variety of additional video output formats, overlay, underlay and keying features, as well as multiple video input formats and signal conditioning options.

The M595, a dual-head display XMC, transfers graphics and video to the host system via a high-speed, eight-port PCIe link. Interfaces include two RGBHV (CRT) channels, an HDTV/TV out port, an LVDS channel and four single-link DVI/HDMI/DP channels through the E4690.

The FPGA, which accepts video via the LVDS and converts it to different formats, supplements the native interfaces with a second HDTV/TV channel, an SDI output port and two STANAG 3350 RGBHV/RGsB channels.

In addition to the capabilities of the M595, the M597, a dual-head graphics and video PMC, adds advanced video functions with analog and digital video input and overlay for exceptional flexibility via 100 MHz connectivity using a 64-bit PCI-X bus. While graphics are being generated in the GPU and overlaid onto one video input for display, a second independent overlay process can run for the second input stream and be delivered on the unique SDI output.

FEATURES

M595 E4690 Graphics XMC

› Single-width XMC
› PCIe 2.0 x8 (8 Lanes) Supporting x1, x2, x4 and x8 Connectivity
› E4690 Embedded GPU @ up to 600 MHz
› Dual Independent Graphic Heads
› 512 MB DDR3 SDRAM @ up to 700 MHz with 128-bit Interface
› 14 Digital Output Channels and SMPTE 259M/292M Serial Output Channel
› 6 Analog Output Channels
› Video Output Routed to XMC Front Panel and P6 I/O Connector IEEE 1386-2001 (Air Cooled) or VITA 20-2001 (Conduction Cooled)

M597 E4690 Graphics and Video PMC

› Single-width PMC
› PCI-X 64-bit @ 133 MHz Connectivity
› E4690 Embedded GPU @ up to 600 MHz
› Dual Independent Graphics Heads
› 512 MB GDDR3 SDRAM @ Up to700 MHz with 128-bit Interface
› 6 Output Channels
› 4 Input Channels and Up to Eight Composite Video Inputs Supporting RS-170A/NTSC/PAL
› Video Capture and Video Overlay Support
› Video I/O Routed to PMC Front Panel and P4 I/O Connectors
› Full 2D/3D Processing Capabilities
› OpenGL 3.x, OpenGL ES 2.0 Processing
› DirectX 10.1
› UVD (Unified Video Decoder) Supporting H.264, VC-1, MPEG-2 Decoding
› Audio Decoding (for HDMI)
› BIOS Flash Memory

Pinnacle Data Systems, Incorporated
6600 Port Road • Groveport, OH 43125 USA
Tel: +1 (614) 748-1150 • Fax: +1 (614) 748-1209
www.pinnacle.com

XMC-GBX Quad Gigabit Ethernet Adaptor

This quad Gigabit Ethernet XMC is a high-performance, low-latency network adaptor providing four high-speed Ethernet interfaces for use with VITA 42.3-compatible VME, PCI Express, CompactPCI®, and AdvancedTCA® processor boards. It is available in three configurations, offering a mix of front and rear port access.

Wide internal data paths eliminate performance bottlenecks. The parallel and pipelined logic architecture is optimized for Gigabit Ethernet and efficiently handles packets with minimum latency. Using widely accepted Intel 82571EB Ethernet Controllers, this adaptor offers up to four 10BASE-T/100BASE-Tx/1000BASE-T copper ports with front-mounted RJ-45 connectors and full-status indicators. Alternatively, up to four SERDES ports are accessible through the Pn4 connector for use via an appropriate copper or fiber-based rear transition module.

FEATURES

› Quad Gigabit Ethernet interfaces – Copper or SERDES
› Up to 4 10BASE-T/100BASE-Tx/1000BASE-T ports with RJ-45 front connectors with status indicators
› Up to 4 rear-accessible SERDES ports via Pn4
› Low-latency data handling
› Efficient packet prioritization
› Enables use of jumbo frames
› Maximum system performance and throughput
› Windows, Linux, and Solaris x86 support
› VITA XMC-compliant interfaces for high bandwidth
› Customization welcomed; extended availability assured

For more information, contact: info.sales@pinnacle.com

www.mil-embedded.com/p45800

Pinnacle Data Systems, Incorporated
6600 Port Road • Groveport, OH 43125 USA
Tel: +1 (614) 748-1150 • Fax: +1 (614) 748-1209
www.pinnacle.com

XMC-E24D/PMC-E24D Dual Display Graphics Module

PDSi offers these high-performance, dual-display graphics modules in both XMC and PMC form factors. Using the ATI Radeon™ E2400 graphics controller from AMD, these modules enable VME, cPCI, and AdvancedTCA systems to take full advantage of AMD's embedded advanced graphics technology. They provide simultaneous independent support of either one digital DVI and one VGA analog display or two VGA displays at 32-bit color and up to 2048 x 1536 resolution.

This module provides the high performance, low power, flexibility, and long life-cycle availability required by many real-world embedded applications in industries such as military/aerospace, industrial control and instrumentation, telecom/datacom, and medical imaging.

FEATURES

› Based on ATI Radeon E2400 graphics processor
› Superior 2D and 3D graphics acceleration
› On-chip GDDR3 video memory
› Dual independent high-performance display interfaces
› DVI-I and analog VGA (full size connectors)
› Dual integrated triple 10-bit DACs for dual RGB output
› Supports analog displays up to QXGA (2048 x 1536)
› 32-bit color depth
› Low-power 65nm design
› Customization and third-party integration welcomed; extended availability assured

For more information, contact: info.sales@pinnacle.com

www.mil-embedded.com/p46816

Innovative Integration

2390-A Ward Avenue • Simi Valley, CA 93065 USA
www.innovative-dsp.com

X6-GSPS

The X6-GSPS integrates high-speed digitizing with signal processing on a PMC/XMC I/O module for demanding DSP applications. The tight coupling of the digitizing to the Virtex-6 FPGA core realizes architectures for SDR, RADAR, and LIDAR front end sensor digitizing and processing. The PCI Express system interface sustains transfer rates over 2 GB/s for data recording and integration as part of a high performance real-time system. Available with Xilinx Virtex-6 SX315T/SX475T or LX240T.

The X6-GSPS features two, 12-bit 1.8 GSPS A/Ds that can be interleaved to use as one 3.6 GSPS digitizer. Analog input bandwidth of over 2.5 GHz supports wideband applications and undersampling. The sample clock is from either a low-jitter PLL or external input. Multiple cards can be synchronized for sampling and down-conversion.

FEATURES

> Two 1.8 GSPS, 12-bit A/D channels
> Single channel interleaved @ 3.6 GHz
> ±1V, AC-Coupled, 50 ohm, SMA inputs
> 4 Banks of 1 GB DRAM (4 GB total)
> Ultra-low-jitter programmable clock
> Gen2 x8 PCI Express providing 2 GB/s sustained transfer rates
> PCI 32-bit, 66 MHz with P4 to Host card
> 20-25W typical Conduction Cooling per VITA 20
> Ruggedization Levels for Wide Temperature Operation
> Adapters for VPX, CompactPCI, desktop PCI and cabled PCI Express systems

For more information, contact: sales@innovative-dsp.com

www.mil-embedded.com/p52898

Concurrent Technologies Inc

6 Tower Office Park • Woburn, MA, 01801 USA
781 933 5900
www.gocct.com

VX 81x/09x – 2nd Generation Intel Core Processor – VME/VXS

The VX 81x/09x is a PC-compatible high performance VME or VXS processor board supporting a choice of 2nd Generation Intel Core processors and up to 16 GB of DDR3 ECC SDRAM. This single-slot board features 2 PMC/XMC sites, a variety of I/O interfaces and optional on-board mass storage. The board supports VITA 41.4 dual x4 PCI Express backplane fabrics and VITA 41.6 Gigabit IEEE 802.3 (1000BASE-BX) ports. In addition to the commercial-grade product, the board is also available in extended temperature and rugged air and rugged conduction cooled variants.

The high processing performance, the wealth of I/O interfaces, and the choice of environmental specifications ensure that this board is suitable for a wide range of applications, in particular within the defense, security and aerospace markets.

To ease integration, many of today's leading operating systems, including Windows, Linux, and VxWorks, are supported.

FEATURES

> Choice of 2- or 4-core 2nd Generation Intel Core Processor
> Up to 16 GB of DDR3-1333 ECC SDRAM
> 2 x SATA interfaces plus optional on-board drive
> On-board CompactFlash site
> 2 x PMC/XMC sites
> Dual independent displays
> 2 x serial channels; 6 x USB interfaces
> 2 x 1000 Mbps Ethernet channels
> VME64/320 interface
> Optional VXS P0 connector (VITA 41.4; VITA 41.6)
> Optional BIT support

For more information, contact: info@gocct.com

www.mil-embedded.com/p47634

Mercury Computer Systems, Inc.

201 Riverneck Road • Chelmsford, MA 01824 USA
866-627-6951
www.mc.com

Ensemble IOM-200 XMC

The IOM-200 I/O Mezzanine XMC Module and IOR-280 RTM from Mercury Computer Systems provide the industry's leading 10Gbps Ethernet I/O density. When two IOM-200s are configured with the IOR-280 Rear Transition Module (RTM), up to eight channels of 10Gbps Ethernet can be supported per 6U OpenVPX slot.

The IOM-200 module includes a high-performance Altera Stratix IV GX field programmable gate array (FPGA) all in a single-wide XMC form factor.

FEATURES

› Up to eight 10GbE channels per 6U OpenVPX slot

› Rear Transition Module available supports two IOM-200 XMCs

› Includes high-performance user programmable FPGA

› Mercury-provided Intel-based drivers supported under Linux industry-leading Quad 10Gbps Ethernet I/O density

› Leverages open standards to ensure interoperability

For more information, contact: webinfo@mc.com   www.mil-embedded.com/p52885

Elma Electronic

44350 Grimmer Blvd. • Fremont, CA 94538 USA
510.656.3400
www.elma.com

VPX Handles and Front Panels

Elma's popular handles are utilized by the leading board vendors and OEMs in the industry. Our handle/panel solutions are now available for VPX, including spacing, ergonomics, and EMC.

The Elma handles feature a rugged design that is ergonomic, reliable, and allows easy assembly with just one screw. Easy extraction is achieved through our two-step latching process with push-button release.

Elma's VPX panels are front anodized, rear alodined, and specially designed to accommodate VPX spacing and board and panel assembly. Elma offers a variety of front panel customization services including silk-screening, digital printing, and custom cutouts.

Contact Elma for a sample today!

FEATURES

› IEEE and low-profile handles, ideal for VPX, VXS, VME64x, cPCI and custom

› Hot swap and non-hot swap handle versions

› Rugged and ergonomic design

› VPX panel sizes of 1.0" and .80"

› Panels accommodate offset spacing for PCBs

› Precision milling for exact panel cutouts

› Front panel customization specialists

› EMC protection

For more information, contact: sales@elma.com   www.mil-embedded.com/p53129

Hypertronics

16 Brent Drive • Hudson, MA 01749 USA
978-568-0451
www.hypertronics.com

KVPX Featuring Hypertac Technology

The KVPX connector is a high-density, high-speed modular inter-connect system optimized for differential pair architectures on a 1.8mm x 1.35mm grid, and is fully footprint-compatible with VITA 46 and VITA 48 standards. Through the use of the proven Hypertac® hyperboloid contact, KVPX connectors provide immunity to shock and vibration fretting, numerous linear paths of contact and a self-cleaning wipe action that results in consistently better signal integrity for unparalleled performance in all dynamic environmental conditions. KVPX connectors are manufactured with 30 percent glass-filled liquid crystal polymer (LCP) insulators and potted contacts to allow withstanding the high temperatures associated with soldering. The LCP insulators also exceed NASA space requirements for outgassing.

Other benefits the Hypertac® contact technology offers include low insertion/extraction forces, low contact resistance, and high cycle life.
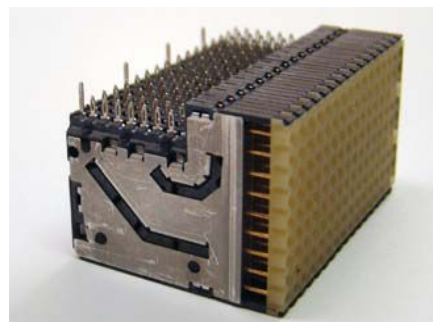
FEATURES

› Fully footprint-compatible with VITA 46 and VITA 48 standards
› Designed for 6.25 Gbps data rate performance scalable to 10 Gbps
› 100 ohm impedance for differential pair configuration
› Differential, single-ended, and power-separable interface offering 63 differential signals/25.4mm and 70 single-ended signals/25.4mm
› Reliable, Hypertac® Hyperboloid contact technology
› ESD protection supports 2-level maintenance designs
› 0.56mm (0.022") diameter via backplane connector
› Flexible, modular design for standard 3U and 6U

For more information, contact: info@hypertronics.com

www.mil-embedded.com/p52500

Kontron

14118 Stowe Drive • Poway, CA 92064-7147 USA
888-294-4558
www.kontron.com

Kontron 3U VPX PCI Express and Ethernet hybrid switch VX3905

The Kontron 3U VPX PCI Express and Ethernet hybrid switch VX3905 is the ideal partner for the centralized backplane to efficiently handle a high bandwidth. It provides up to 24 ports with 32-lane PCI Express Gen 1/Gen 2 switching and additional 9-port Gigabit Ethernet switching capabilities for the control plane. With such a feature set in a single 3U VPX slot, very efficient and compact computer architectures can be designed to satisfy the most stringent SWaP-C requirements in high-performance embedded computing.

The Kontron VX3905 is available in an air-cooled version for ambient temperatures from 0° to +55°C, and in a rugged conduction-cooled version for the extended temperature range from -40° to +85°C.

FEATURES

› Compliant with OpenVPX VITA 65 profile SLT3-SWH-6F6U-14.4.1
› Up to 24-port/32-lane PCIe switch
› 9-port Gigabit Ethernet switch
› Air-cooled and conduction-cooled builds

For more information, contact: info@us.kontron.com

www.mil-embedded.com/p53033

Kontron

14118 Stowe Drive • Poway, CA 92064-7147 USA
888-294-4558
www.kontron.com

Kontron 3U VPX SBC VX3035

Integrating the Intel® Core™ i7-2655LE processor, Intel® HD graphics and features such as Intel® Turbo Boost technology and Intel® Advanced Vector Extensions (Intel® AVX), the Kontron VX3035 defines a new performance class for SWaP (size, weight and power)-optimized VPX applications.

The Kontron VX3035 supports VXFabric™, allowing data flow applications using IP sockets to enjoy high-bandwidth interboard communication at PCIe DMA hardware speed. The Kontron VX3035 is available standard air-cooled, rugged air-cooled or rugged conduction-cooled for extreme environmental conditions.

FEATURES

> Intel® Core™ i7-2655LE
> Up to 8 GB dual-channel 1333 SDRAM with ECC
> Support for Windows Embedded Standard 7, Linux and VxWorks
> Supports Kontron VXFabric™
> Standard air-, rugged air- or rugged conduction-cooled versions
> Available also as turnkey evaluation/development platforms

For more information, contact: info@us.kontron.com   **www.mil-embedded.com/p53034**

COTS Collection: Boards, Carriers, and Mezzanines: VITA 46 VPX

Triple E

358 North Street • Randolph, MA 02368 USA
1-888-444-1644
www.tripleease.com

Triple E VPX – The Complete Ruggedized Packaging Toolbox

Triple E provides the complete mechanical packaging solution for VPX. The Triple E portfolio includes all-metal front panel handles, front panels with ruggedized, integral pb holders, metal card guides, card decks, and card cages – all the components for configuring a robust, reliable system!

The all-new VPX-TE handle features stainless construction with an ergonomic push-button design and sturdy locking feature.

The new VPX card guides and rails feature an integrated stop function and a rugged U channel front rail. The line is available in a wide range of finishes, including powder coat, chromate, and yellow chromate. Just add a backplane and the subrack is complete.

Contact us now for your datasheet on the Triple E VPX Toolbox!

For additional information visit: www.tripleease.com/VPX

FEATURES

> VPX all-metal card guide with integrated stop feature
> Finish options: powder coat, chromate, yellow chromate
> Stainless metal VPX front panel handle with secure push-button locking
> U Channel-extruded front panels with metal gasket (standard) – optional foam-over-fabric gasket
> Integral pb holder option for secure mounting
> U channel front rail for rigidity and enhanced front panel interface

For more information, contact: tripleesales@xtech-outside.com   **www.mil-embedded.com/p53133**

Annapolis Micro Systems, Inc.

190 Admiral Cochrane Drive, Suite 130 • Annapolis, MD 21401 USA
410-841-2514
www.annapmicro.com

WILDSTAR 6 for OpenVPX

Annapolis Micro Systems is a world leader in high-performance, COTS FPGA-based boards and processing for RADAR, SONAR, SIGINT, ELINT, DSP, FFTs, communications, Software-Defined Radio, encryption, image processing, prototyping, text processing, and other processing-intensive applications.

Our 14th-generation WILDSTAR 6 for OpenVPX uses Xilinx's newest Virtex-6 FPGAs for state-of-the-art performance. It accepts one or two I/O mezzanine cards in one VPX slot or up to 4 in a double wide VPX slot, including Single 1.5 GHz 8-bit ADC, Quad 250 MHz 12-bit ADC, Single 2.5 GHz 8-bit ADC, Quad 130 MHz 16-bit ADC, Dual 2.3/1.5 GSps 12-bit DAC, Quad 600 MSps 16-bit DAC, Universal 3Gbit Serial I/O (RocketIO, 10 Gb Ethernet, InfiniBand), and Tri XFP (OS 192, 10G Fibre Channel, 10 Gb Ethernet). Our boards work on a number of operating systems, including Windows, Linux, Solaris, IRIX, ALTIX, and VxWorks. We support our board products with a standardized set of drivers, APIs, and VHDL simulation models.

Develop your application very quickly with our CoreFire™ FPGA Application Builder, which transforms the FPGA development process, making it possible for theoreticians to easily build and test their algorithms on the real hardware that will be used in the field. CoreFire, based on dataflow, automatically generates distributed control fabric between cores.

Our extensive IP and board support libraries contain more than 1,000 cores, including floating point and the world's fastest FFT. CoreFire uses a graphical user interface for design entry, supports hardware-in-the-loop debugging, and provides proven, reusable, high-performance IP modules. WILDSTAR 6 for OpenVPX, with its associated I/O Cards, provides extremely high overall throughput and processing performance. The combination of our COTS hardware and CoreFire allows our customers to make massive improvements in processing speed, while achieving significant savings in size, weight, power, person-hours, dollars, and calendar time to deployment.

Annapolis is famous for the high quality of our products and for our unparalleled dedication to ensuring that the customers' applications succeed. We offer training and exceptional special application development support, as well as more conventional support.
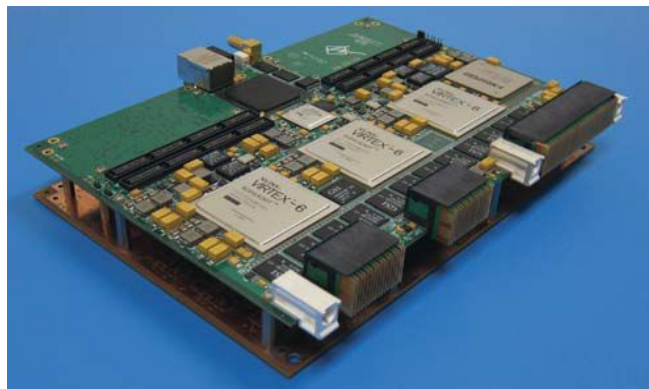
FEATURES

› Up to three Virtex-6 FPGA processing elements – XC6LX240T, XC6LX365T, XC6LX550T, XC6SX315, or XC6SX475

› Up to 7 GB DDR2 DRAM in 14 banks or up to 448 MB DDRII or QDRII SRAM

› OpenVPX backplane

› 80 x 80 crossbar connecting FPGAs and VPX backplane

› 1 GHz 460EX PowerPC onboard host

› 4X PCIe controller

› Programmable Flash to store FPGA images and for PCI controller

› Full CoreFire Board Support Package for fast, easy application development

› VHDL model, including source code for hardware interfaces and ChipScope access

› Host software: Windows, Linux, VxWorks, etc.

› Available in both commercial and industrial temperature grades

› Proactive thermal management system – board-level current measurement and FPGA temperature monitor, accessible through Host API

› Save time and effort and reduce risk with COTS boards and software

› Achieve world-class performance – WILD solutions outperform the competition

› Includes one-year hardware warranty, software updates, and customer support; training available

BittWare

9 Hills Avenue • Concord, NH 03301 USA
603-226-0404
www.bittware.com

BittWare SIGNAL PROCESSING SYSTEMS

Commercial & Rugged Altera Stratix® IV GX 3U VPX Board

BittWare's S4-3U-VPX (S43X) is a commercial or rugged 3U VPX card based on the high-density, low-power Altera Stratix IV GX FPGA. The Stratix IV GX is designed specifically for serial I/O-based applications, creating a completely flexible, reconfigurable VPX board.

BittWare's ATLANTiS FrameWork and the FINe Host/Control Bridge greatly simplify application development and integration of this powerful board.

The board provides a configurable 25-port SerDes interface, supporting a variety of protocols, including Serial RapidIO, PCI Express, and 10 GigE. The board also features 10/100/1000 Ethernet, and up to 4 GB of DDR3 SDRAM. Providing enhanced flexibility is the VITA 57-compliant FMC site, which supports 10 SerDes, 60 LVDS pairs, and 6 clocks.

FEATURES

› VITA 57 FMC site with options available for high-speed data conversion, providing up to 8-bit 2.5 GSPS ADC, and up to 16-bit 1 GSPS DAC

› High density Altera Stratix IV GX supported by BittWare ATLANTiS™ FrameWork for FPGAs

› BittWare FINe™ Host/Control Bridge provides control plane processing and interface

› Fully connected to VPX: GigE, 15 SerDes, 32 LVDS

› Additional I/O: 10/100 Ethernet, RS-232, JTAG

For more information, contact: info@bittware.com

www.mil-embedded.com/p52808

CES – Creative Electronic Systems

Avenue Eugene-Lance 38 • 1212 Grand Lancy, Geneva Switzerland
+41.22.884.51.00
www.ces.ch

CES CREATIVE ELECTRONIC SYSTEMS

ETS-8227

The ETS-8227 is the world's first 3U, OpenVPX multi-fabric switch board featuring PCIe, SRIO and GbE switches for ground or airborne applications. This board allows switching between PCI Express (Gen1/Gen2) and Serial RapidIO (1.3/2.1) data planes. In addition, the two onboard crosspoint switches permit flexibility of the payload profile configuration in accordance to OpenVPX. The ETS-8227 also incorporates a Gigabit Ethernet switch for the control plane that includes Layer 2 switching and aggregation features. An integrated advanced board management controller provides high-speed system status monitoring, logging and dynamic reload functions in real time.

FEATURES

› 3U OpenVPX form factor (VITA 65)

› One PCI Express Gen2 switch (12 ports) for data plane

› One Serial Rapid IO 2.1 switch (12 ports) for data plane

› One Gigabit Ethernet switch (8 ports) for control plane

› Two crosspoint switches (40x40) for programmable payload profile configuration (PCIe x4 Gen2, SRIO x4, GbE)

› Integrated advanced board management controller

› Rear I/O transition module for additional interfaces

For more information, contact: ces@ces.ch

www.mil-embedded.com/p52941

Innovative Integration

2390-A Ward Avenue • Simi Valley, CA 93065 USA
www.innovative-dsp.com

VPX-COMEX

The VPX-COMEX is a 3U OpenVPX system controller CPU card that integrates an Intel CPU COM Express CPU module with SRIO switch, system timing and triggering features, and a Spartan-6 FPGA core. The VPX-COMEX supports VPX systems with up to 4 other PCIe/SRIO cards and timing/triggering support for embedded communications, instrumentation, and data acquisition applications.

The CPU core is a COM Express module, a PICMG industry standard, that is an Intel Architecture i7 or i5 CPU with QM57 chipset. The COM Express module runs Windows, Linux, or VxWorks, providing a familiar and easy-to-use software environment that is abundant in tools and applications, resulting in dramatically lower time-to-market than other CPU architectures.

The VPX-COMEX supports 4 expansion cards using PCI Express and SRIO. The PCI Express root complex has x4, x4, x1, and x1 Gen1 endpoint support, provided by the chipset. The SRIO data plane complements the PCIe connectivity with a central switch, providing 8 SRIO Gen2 links (5 Gbps each).

FEATURES

> COM Express Type 6 CPU module
> – Intel i7 @ 2.53 GHz/i5 @ 2.4 GHz
> – Up to 8 GB 1066 MHz DDR3 ECC memory
> PCI Express Root supports x4/x4/x1/x1 Gen1 endpoints
> GbE, USB 2.0, 3x SATA300, DisplayPort
> Integrated 1.8" SATA SSD up to 256 GB
> Serial RapidIO Data Plane Switch
> – 8x SRIO Gen2 (5 Gbps) ports to VPX
> – Front panel QSFP x4 optical link
> Precision timing/triggering support
> – PLL with 125 KHz to 1 GHz tuning range with -110 dB phase noise @ 10 KHz
> – 10 MHz, 0.28 PPM clock reference
> – Reference input support for GPS integration
> – Synchronized triggering outputs
> Xilinx Spartan-6 LX45T to LX120T FPGA core
> < 50 W typical; conduction or forced-air cooling
> Operating Environment: 0°C to 60°C, 0 to 100% RH, 30g shock, 0.04 g^2/Hz random vibe

For more information, contact: sales@innovative-dsp.com   **www.mil-embedded.com/p52900**

Innovative Integration

2390-A Ward Avenue • Simi Valley, CA 93065 USA
www.innovative-dsp.com

VPX6-COP

The VPX6-COP is a flexible FPGA coprocessor card that integrates a Virtex-6 FPGA computing core with an industry-standard FMC I/O module on a 3U OpenVPX card.

The FPGA computing core features the Xilinx Virtex-6 FPGA family in densities up to LX550 and SX475. The SX475 provides over 2000 DSP MAC elements operating at up to 500 MHz. The FPGA core has two 9 MB QDRII + SRAM banks, two 256 MB LPDDR2 DRAM banks, and a 128 MB DDR3 bank. Each memory is directly connected to the FPGA and is fully independent.

For system communications the VPX6-COP has a PCI Express and two SRIO/Aurora interfaces. The PCIe port is a x8, Gen2 interface capable of up to 2 GB/s sustained operation with 4 GB/s burst rate. Two additional x4 system ports support either SRIO, Aurora, or custom protocols.

FEATURES

> 3U OpenVPX FPGA coprocessor card
> FMC I/O site (VITA 57) with 8x 5 Gbps MGT lanes, 80 LVDS pairs (LA, HA, HB full support)
> 2 Banks of 256 MB DRAM (512 MB total)
> 2 banks of 9 MB QDRII + SRAM (18 MB total)
> 128MB DDR3 DRAM
> VPXI system-timing features supporting global and local timing and triggering features
> Gen2 x8 PCI Express providing 4 GB/s burst and 2 GB/s sustained transfer rates
> Two Serial RapidIO or Aurora ports supporting x4 Gen2 (2 GB/s)
> < 15 W typical, excluding FMC
> Ruggedization levels up to L4 forced air or conduction cooling
> 40 g shock, 9 g sine vibration, and 0.1 g^2/Hz random vibe

For more information, contact: sales@innovative-dsp.com   **www.mil-embedded.com/p52896**

Innovative Integration

2390-A Ward Avenue • Simi Valley, CA 93065 USA
www.innovative-dsp.com

VPXI-ePC

The VPXI-ePC is a 3U VPX embedded computer system that provides a performance architecture for instrumentation, signal processing and embedded computing applications. Four expansion slots plus an I/O/HDD drive are in a compact, half-rack 4U enclosure. The VPX-COMEX CPU combines an Intel Architecture COM Express CPU module with timing and communications features. Real-time, data-intensive operation is built on multiple, high-performance data planes, employing both centralized and mesh topology inter-connections between the cards.

VPXI integrates timing features into the VPX architecture, providing synchronized high-performance clock and trigger features to each slot. Peripheral slots receive a dedicated clock and trigger input, as well as several shared coordination signals. These signals are used by VPX-COP, X6 and X3 I/O card families and support simultaneous and coordinated sampling. An optional high-precision GPS can be used with the VPX-COMEX as a timing reference.

The backbone of the system in the VPXI-ePC comprises the PCIe and SRIO planes in the VPXI system. Each plane can sustain data rates over 500 MB/s concurrently on PCIe and SRIO planes. The SRIO data plane also has a x4 optical link on the front panel supporting 12 Gbps connection.

The VPXI may host one VPX-COMEX CPU card to create a Windows/Linux/VxWorks-compatible PC. The VPXI-ePC can run the same applications as a desktop system. Performance OSs such as Linux, Xenomai and VxWorks are available for real-time applications.

VPXI-ePC

FEATURES

› VPX for Instrumentation
  – 3U OpenVPX embedded computer system
  – Integrated timing and triggering
  – Advanced multiple-plane connectivity
  – Rugged with wide-temperature options

› Embedded PC
  – Runs Windows, Linux or VxWorks
  – COM Express module with Intel i5/i7 CPU with up to 8 GB memory
  – Gb Ethernet, 4x USB, DisplayPort video
  – 256 GB SSD + up to 3 removable drives

› Half (1/2) rack, 4U system
  – Six slots total: CPU slot + 4 OpenVPX Peripheral Slots + Storage/I/O Slot
  – Compatible with many OpenVPX cards
  – Up to 3 HDD/SSD storage slots
  – Supports Innovative X3/X5/X6 and VPX-COP

› Integrated timing and triggering features
  – Synchronized multi-card sampling
  – Internal or external clock/references
  – Generates low phase noise sample clocks from 0.125 to 1 GHz
  – 10 MHz, 0.5 PPM stable clock reference
  – Optional GPS-disciplined reference

› Advanced architecture supports multiple data planes
  – PCI Express and SRIO planes
  – Mesh interconnects for all I/O cards
  – Front panel x4 optical link for SRIO IO

› Rear Terminal Modules for I/O and CPU slots
  – Forced air cooling with upper and lower fans
  – Integrated 500W power supply
  – Expands to additional VPXI chassis using cable PCI Express option

APPLICATIONS

› Remote, Autonomous I/O
› Mobile Instrumentation
› Distributed Data Acquisition
› Signal Processing Clusters

For more information, contact: sales@innovative-dsp.com

www.mil-embedded.com/p52899

Mercury Computer Systems, Inc.

201 Riverneck Road • Chelmsford, MA 01824 USA
866-627-6951
www.mc.com

Echotek Series SCFE-V6-OVPX

The Echotek® Series SCFE-V6-OVPX Virtex-6 FPGA Processing Engine from Mercury Computer Systems provides high-performance processing for applications requiring extreme FPGA processing power, such as EW, ELINT, SIGINT, RADAR, Commercial Wireless, and Deep Packet Inspection. Utilizing top-end FPGAs available from market leader Xilinx®, the SCFE-V6-OVPX solves the toughest signal processing problems in a cost-effective form factor. VITA 57 FPGA Mezzanine Card (FMC) mezzanine sites and OpenVPX Rear Transition Modules (RTMs) provide diverse sensor and I/O entry points, both analog and digital. When incorporated with the Mercury Ensemble™ multi-computing OpenVPX modules, the SCFE-V6-OVPX represents an essential building block in powerful heterogeneous subsystems that can be tailored to specific applications and platforms.

FEATURES

› Next-generation architecture for high-end defense and commercial applications such as SIGINT, EW, Communications, RADAR, Commercial Wireless, and Deep Packet Inspection
› Ultimate processing power with three Xilinx® Virtex™-6 LX240Ts or SX315Ts
› VITA 57 FMC sites for flexible I/O enhancements, including Echotek® Series multi-channel digitizer FMCs
› High-speed OpenVPX interface fully compatible with Ensemble OpenVPX modules
› EchoCore™ IP Advantage
› Flexible Control-Plane Processor
› Flexible I/O with FMC Sites

For more information, contact: webinfo@mc.com   **www.mil-embedded.com/p52893**

Mercury Computer Systems, Inc.

201 Riverneck Road • Chelmsford, MA 01824 USA
866-627-6951
www.mc.com

Ensemble GSC6200

The Ensemble™ 6000 series 6U OpenVPX™ GSC6200 GPU Processing Module from Mercury Computer Systems, Inc. harnesses the tremendous compute power of graphics processing units (GPUs) for rugged, high-performance, embedded signal and image processing in a broad range of defense and commercial applications, including radar signal processing, electro-optical and infrared image processing, electronic warfare and other applications requiring intense computational capabilities on large streams of data. Apply GPUs to accelerate compute-intensive applications including Fast Fourier Transforms (FFTs), Constant False Alarm Rate (CFAR), QR Decomposition (QRD) Synthetic Aperture Radar (SAR). The GSC6200 uses 2 NVIDIA GPUs based on their latest generation Fermi architecture, which provides significant enhancements over the previous generation GPUs, including a larger number of parallel processing cores, larger memory capacity and higher memory bandwidth using GDDR5.

FEATURES

› Dual MXM sites support 1 NVIDIA GeForce GTX 460M Fermi Architecture GPU, each with two single-link DVI interfaces
› 48-lane, 12-port PCI Express Gen2 switch to backplane
› Software-controlled management features
› CUDA software environment for NVIDIA-based GPUs
› Supports up to 2 high-end NVIDIA® compute nodes per 6U slot
› MXM form factor preserves rapid technical insertion of latest GPU architectures
› Advanced power management capabilities
› Conduction-cooled and air-cooled versions

For more information, contact: webinfo@mc.com   **www.mil-embedded.com/p52886**

Mercury Computer Systems, Inc.

201 Riverneck Road • Chelmsford, MA 01824 USA
866-627-6951
www.mc.com

Ensemble SFM3010

The Ensemble™ 3000 Series SFM3010 3U OpenVPX™ Switch Fabric Module from Mercury Computer Systems combines Serial RapidIO® data-plane switching with a managed Gigabit Ethernet control-plane switch for maximum scalability. The SFM3010 also implements the OpenVPX Chassis Manager function for mission-critical system management. With the module's XMC mezzanine site, users can stream I/O directly onto the data plane, delivering data to any processing element in the system without significant overhead. By combining these functions on a single 3U OpenVPX module, the SFM3010 expands subsystem capabilities for size, weight, and power (SWaP)-constrained environments.

FEATURES

> Control plane communications supported by a Broadcom Gigabit Ethernet switch with integrated MIPS processor core

> Data plane communications supported by a Tundra Serial RapidIO switch with six 4x links to the backplane and two 4x links to the J15 XMC connector

> XMC site supported by Serial RapidIO per VITA 42.2

> IPMI controller for system management functions

> Air-cooled and conduction-cooled models available

For more information, contact: webinfo@mc.com

www.mil-embedded.com/p52883

Mercury Computer Systems, Inc.

201 Riverneck Road • Chelmsford, MA 01824 USA
866-627-6951
www.mc.com

Ensemble HCD3210

The Ensemble™ 3000 Series HCD3210 3U OpenVPX™ Processing Module from Mercury Computer Systems is a single-slot solution for data acquisition and processing. By combining the performance logic of a highly capable Xilinx® Virtex™-6 FPGA with the real-time processing power of a Freescale™ MPC8640D dual-core processor, the HCD3210 supports multi-stage processing for many applications. The module's XMC mezzanine site delivers I/O to these powerful processing elements, supporting high-performance processing for size, weight, and power (SWaP)-constrained platforms. The multi-plane OpenVPX architecture allows users to scale, knitting together multiple HCD3210s into a powerful subsystem.

FEATURES

> General-purpose processing via dual-core MPC8640D processor running at 1.06 GHz, supported 1 GB of DDR2 SDRAM

> Xilinx® Virtex™-6 FPGA supported by two 9 MB banks of QDRII SRAM and 128 MB of DDR3-900M SDRAM

> High-bandwidth communications, including native support for RapidIO and PCI Express® (PCIe)

> XMC site supported by PCI Express (PCIe) per VITA 42.3

> IPMI controller for system management functions

> MultiCore Plus® software environment with MC SAL

For more information, contact: webinfo@mc.com

www.mil-embedded.com/p52884

Mercury Computer Systems, Inc.
201 Riverneck Road • Chelmsford, MA 01824 USA
866-627-6951
www.mc.com

MERCURY
COMPUTER SYSTEMS ™

Ensemble LDS6521

The Ensemble™ 6000 Series OpenVPX Intel® Core i7 Quad-Core Next Generation LDS6521 Module combines a powerful Sandy Bridge mobile-class quad-core Intel® 2nd Generation Core i7 processor, a high-performance FPGA for both fabric bridging and user-application functions, and high-bandwidth on-board and off-board communication fabrics in a single 6U OpenVPX slot.

The LDS6521 low-density server provides a next-generation architecture that balances the disruptive computational capabilities of the Intel 2nd Generation Core i7 processor with key high-bandwidth I/O interfaces, providing a powerful and scalable computing architecture that is well aligned with high-end radar, electronic warfare, and image processing applications.

FEATURES

› One 10/100/1000BASE-T GigE connection: front panel on air-cooled, backplane on conduction-cooled
› One additional 10/100/1000BASE-T GigE connection to the backplane
› Two 1000BASE-BX SERDES Ethernet connections to the backplane
› One front panel USB 2.0 interface on air-cooled
› Two backplane USB 2.0 interfaces on both air-cooled and conduction-cooled
› One front-panel eSATA interface provided on air-cooled
› Two SATA interfaces to the backplane
› Eight GPIO lines

For more information, contact: webinfo@mc.com

www.mil-embedded.com/p52892

COTS Collection: Boards, Carriers, and Mezzanines: VITA 65 OpenVPX

Emerson Network Power
2900 S. Diablo Way, Suite 190 • Tempe, AZ 85282 USA
1-800-759-1107 or +1-602-438-5720
www.Emerson.com/EmbeddedComputing

EMERSON
Network Power

3U & 6U VITA 46 VPX & VITA 65 OpenVPX Processor Board

The 6U iVPX7220 and 3U iVPX7223 VITA 46 VPX & VITA 65 OpenVPX™ processor boards feature the dual- & quad-core 2nd gen Intel® Core™ i7 processors with integrated graphics & memory controller and the Mobile Intel® QM67 Express chipset with leading edge I/O functionality. This high compute density platform offers both high speed fabric connectivity with PCI Express and Gigabit Ethernet control plane connectivity with data transfer rates up to 5Gbps.

On-board memory includes up to 16GB DDR3-1333 memory, embedded USB flash and 256KB non-volatile F-RAM. Additional connectivity includes a variety of USB 2.0, serial and SATA ports, GPIO, DisplayPort, VGA and XMC sites for maximum flexibility. An optional 2.5" SATA SSD is also available.

The boards are fully rugged for extreme environments with extended shock, vibration, temperatures and conduction cooling. Software support includes solid and stable BIOS with password protection and a wide range of operating systems.
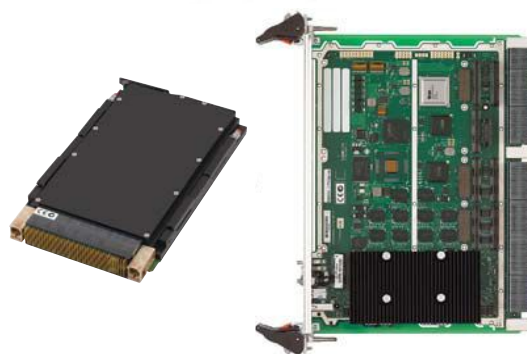
FEATURES

› 2nd generation Intel® Core™ i7 2.20 GHz dual- or quad-core integrated processor & Mobile Intel® QM67 PCH
› Up to 8GB (3U) or 16GB (6U) ECC-protected DDR3-1333 soldered
› VITA 48 REDI two-level maintenance (2LM)
› Extended temperature -40 °C to +85 °C 3U; -40 °C to +71 °C 6U and rugged variants
› Air and conduction cooled
› Designed for a range of industrial, communication and military/aerospace applications

For more information, contact: EmbeddedComputingSales@Emerson.com

www.mil-embedded.com/p53127

Kontron

14118 Stowe Drive • Poway, CA 92064-7147 USA
888-294-4558
www.kontron.com

Kontron 6U VME SBC VM6050

The Kontron VM6050 with the Intel® Core™ i7 processor brings increased performance and reduced development time to new and existing VME system designs.

With 100% I/O compatibility with the Kontron VM6250, its PowerPC sibling, both featuring a VITA 57 FMC interface, the Kontron VM6050 is the ideal bridge between all existing VME designs and modern CPU and I/O performance and price, regardless of the software legacy. It combines extremely high x86 computing and graphics performance with flexible and modular expansion possibilities in four different ruggedization levels.

With Kontron's long-term supply program with availability of 10 years or more, OEMs can further optimize the life cycle and total cost of ownership (TCO) for applications.

FEATURES

› 6U VME high-performance Intel® Core™ i7 processor
› Outstanding performance in VME format with SSE 4.2 CPU instruction set
› Exceptional I/O versatility
› Commercial and rugged versions
› Extended life cycle with adapted long-term support program

For more information, contact: info@us.kontron.com   **www.mil-embedded.com/p53036**

New VITA Standards: VITA 65 OpenVPX

CSP Inc. MultiComputer Division

43 Manning Road • Billerica, MA 01821 USA
978-663-7598
www.cspi.com/multicomputer/

3220Q

The 3220Q 6U OpenVPX Blade Server, with a dual socket Intel® Xeon® processor and QuickPath Interconnect (QPI), delivers superior performance-per-watt and low latency while operating in a standards-based OpenVPX environment.

Offering an open software stack, the 3000 SERIES OpenVPX blades support MPI over InfiniBand, cluster management tools, a Linux 2.6 SMP-based distribution or VxWorks. Continued support for the same message passing protocol between processes allows existing customers to migrate without any change at the Application Layer. This strategy provides customers with an upgrade path to the most advanced technologies while offering widespread interoperability and enabling software reuse. Likewise, the software development environment preserves support for Wind River's Workbench open device software development suite, and CSPI's optimized signal processing libraries.

FEATURES

› Dual Socket Intel® Xeon® Processors
  – 4 Cores/8 Threads per socket
  – 8 MB On-chip Cache per socket
  – 2.13 GHz Core clock
  – 4.8 Gbps QuickPath Interconnect (QPI) between sockets
  – 8/16 GB DDR3 ECC-protected device-down Memory
  – 136 GFLOPS Peak Performance
› 32 Lanes of PCIe Gen2 to the OpenVPX Data and Expansion Planes
› Two 1000BASE-T Interfaces to the VITA-46.6 Ethernet Control Plane
› Full IPMC Implementation for compliance with future VITA 46.11 Management Standard

For more information, contact: info@cspi.com or 1-800-325-3110   **www.mil-embedded.com/p52811**

Military Embedded Systems   Resource Guide

WinSystems, Inc.

715 Stadium Drive • Arlington, TX 76011 USA
817-274-7553
www.WinSystems.com

Extended Temperature CompactFlash

These CompactFlash cards are targeted for applications that need industrial-grade reliability, industry-standard compatibility, and IDE hard disk drive emulation for fast program and data storage. Our CompactFlash cards have greater than 2 million program/erase cycles and unlimited reads while maintaining fast transfer speeds of up to 66MB/s in burst mode.

Fully operational from -40°C to +85°C, these CompactFlash cards are available in storage densities from 128MB to 16GB, and fit any computer, SBC, or instrument with a CompactFlash socket. Since they are True IDE Mode and ATA-3 compliant, they are compatible with different operating systems such as Linux, Windows®, and other RTOSs without requiring special drivers.

We offer in-stock product availability plus free technical telephone support with a factory applications engineer.

www.IndustrialCompactFlash.com

FEATURES

› 128MB to 16GB storage capacity
› Up to 66MB/s burst with 37MB/s read and 16MB/s write sustained
› Sophisticated error-checking and wear-leveling algorithms
› Withstands 2,000Gs shock and 16.3Gs vibration
› -40°C to +85°C operating temperature
› In-stock availability and RoHS-compliant

For more information, contact: Info@WinSystems.com

www.mil-embedded.com/p35888

Microsemi Corporation – Power & Microelectronics Group

3601 E. University • Phoenix, AZ 85034 USA
1-888-SSD-4-MIL (1-888-773-4645)
www.microsemi.com/pmgp

TRRUST-Stor™

A New Level of SSD Security, Designed for Defense

With encryption, unparalleled ruggedization, and blazing fast erase, Microsemi's TRRUST-Stor™ is the first and only solid state drive designed from the ground up for aerospace, defense, and military applications. The environmentally ruggedized, high-reliability TRRUST-Stor is the industry's first SSD capable of withstanding zero-failure testing at vibration levels up to 42 Grms, making it ideal for the most severe defense application environments. TRRUST-Stor delivers security, reliability, and performance unmatched by current commercial SSD offerings for mission-critical applications such as vetronics, UAVs, surveillance, data recorders, digital map storage, avionics, GPS/communications systems, and ruggedized field computers.

Why trust your data to anything less?

FEATURES

› 50 or 100 GB densities
› Hardware-implemented AES-256 encryption
› Tamper-resistant features available
› EOL management
› TRRUST-Purge™ renders data irrecoverable in milliseconds
› Full drive erase in less than 4 seconds
› Hardware-based authentication
› Low power operation
› MTBF greater than 2,000,000 hours
› Military & government agency sanitization protocols
› Built-In-Self-Test (BIST)

For more information, contact: jack.bogdanski@microsemi.com

www.mil-embedded.com/p53130

InnoDisk USA Corporation

43130 Osgood Road • Fremont, CA 94539 USA
T: (510) 314-8393 • F: (510) 314-8394
www.innodisk.com

InnoRobust Ultra SSD

InnoDisk Corporation, a DRAM and Flash Storage Solution Designer and Manufacturer, introduces the InnoRobust Series, designed specifically for Military and Mission-Critical Applications. When secure information needs to be stored, HDD is not an option – InnoRobust SSD is the answer!

InnoRobust with Write Protect, Destroy, and Secure Erase supports the latest requirements demanded by the US Armed Forces. The InnoRobust Series is designed for RAID Applications, Rugged Servers, Rugged Storage Unit, Submarines, Air, Sea, and Land Vehicles, and UAVs.

InnoRobust Specifications:

Performance:
Sustained Read/Write (MB/s):	175/90
Random Read/Write (IOPS@4KB):	2400/150

Power Consumption:
Read/Write mA (max.):	740/770
Idle:	460mA (max.)

Environmental:
Altitude:	80,000 feet
Operating Temp.:	Industrial 0 °C to 70 °C, W/T -40 °C to +85 °C
Storage Temp.:	-55 °C to +95 °C
Shock (Operating):	1500G
Vibration (Operating):	20G
Humidity:	5-95%, non-condensing
Write Protect:	Support
Data Retention:	10 years
Dimension:	99.88mm x 69.63mm x 9.3mm (LxWxH)

InnoDisk is the leading manufacturer of industrial storage devices worldwide. The elite R&D team holds many product patents and is well experienced in designing industrial-grade storage devices for embedded systems. InnoDisk provides comprehensive solutions for embedded systems and specialized applications. InnoDisk is among the most reliable and highest performing vendors globally.

For more information on the full line of InnoDisk products, including specifications, please visit our website at www.innodisk.com or contact us for more information.

InnoRobust 1.8" SATA SSD
Capacities: 8GB - 64GB

InnoRobust Ultra 2.5" SATA SSD
Capacities: 16GB - 256GB

FEATURES

› Capacity: 16GB to 256GB
› Single Level Cell (SLC) NAND Flash
› Compact Size: 2.5" SSD and 1.8" SSD
› Support SATA and PATA plug and play
› Hot Swappable
› Data Security/Protection: Quick Erase and Secure Erase, Write Protect/Destroy
› Conformal Coating: Preventing Corrosion/MIL 1-46058C silicon conformal coating
› Low Power Consumption, Operates in mA
› MIL-STD-810F 503.4 Thermal Shock Procedure II; -40 °C to +85 °C, 10 Cycles
› MIL-STD-810F 512.5C-1 Specification, Vibration 20G
› MIL-STD-810F 516.5 Procedure IV Transit Drop Specification, Shock of 1500G, 3 axis for XYZ
› 40g/11ms/18 shocks (3 shocks for = X, Y, Z axis)
› Quick Erase – Ability to erase all data, master boot records, and write over in a known pattern in 35 seconds on a 128GB SSD
› Secure Erase complies with military data elimination standard
› The erase feature can be performed manually or by command
› Once the Erase Command has been engaged there is no stopping the device until the SSD has been fully erased and written over by a standard pattern.
› **InnoRobust complies to the below guidelines:**
  • USA-AF AFFSI 5020
  • DoD 5220.22-M
  • USA Navy NAVSO P-5239; NSA 9-12
  • USA Army 380-19
  • NISPOMSUP Chap. 8, Sect. 8-501

For more information, contact: usasales@innodisk.com

www.mil-embedded.com/p46024

Military Embedded Systems   Resource Guide

Memoright Corporation

5104 Old Ironsides Drive, Suite 113 • Santa Clara, CA 95054 USA
409-961-4007
www.memoright.com

Memoright 2.5" GTR II series SLC SSD

The Latest Ruggedized SSD Innovation Offers Utmost Performance and Reliability for Industrial and Military applications

Memoright, a Taiwan-based company dedicated to technological development and integration for SSDs (solid-state drives) and high-end storage equipment, announces the launch of the second-generation industrial- and military-grade ruggedized GTR II SLC SSD. Memoright brings the best performance ruggedized SSD for military, defense, aviation and industrial applications.

Memoright's self-developed circuit design and close-examination manufacturing process for the GTR II provides high performance and 100% constant speed. GTR II supports SATA II interfaces and is available in 16GB, 32GB, 64GB and 128GB capacities. With its emphasis on stable and reliable performance, Memoright SSD has been widely adopted for military and industrial storage purposes.

Performance:
GTR II is designed specifically for industrial/military storage, with a -40 ˚C ~ +85 ˚C operating temperature, as well as handling extreme altitude, vibration and humidity. GTR II provides outstanding performance (seq. read/write: 210MB/210MB; IOPS read/write: 10,000/500) for military/industrial storage applications, which require the utmost reliability and performance.

Security:
GTR II supports Security Erase and Destroy functions and satisfies all military standards under "MIL-STD-810F", which is the standard test method approved for use by all departments and agencies in the US Department of Defense (DoD).

Reliability:
GTR II also features Memoright's proprietary and unique "In-Drive UPS", which ensures 100% data integrity at all times. This avoids the occurrence of abnormal bad blocks and prolongs the useful life of SSDs.

FEATURES

› Meets both U.S. MIL-STD-810F/G and China GJB, the strictest standards

› Wide Temperature : -40 °C ~ +85 °C

› Seq. Read/Write: 210MB/210MB

› IOPS: 10,000/1,500

› SLC NAND Flash

› Supports SATA II

› Capacity: 16/32/64/128 GB

› Humidity: 5%~95%

› Supports severe conditions in extremely high pressure from -1,000 to 80,000 Feet

› In-Drive UPS: A specific energy storage element guarantees no data loss during any unexpected power spikes, dips or failures

› Security Erase: Memoright provides users both commercial and military Secure Erase modes that can be triggered either by the hardware button or a sequence of commands.

For more information, contact: sales@memoright.com

Mass Storage: Solid State Disk (SSD)

Phoenix International

812 W. Southern Ave. • Orange, CA 92865 USA
714-283-4800
www.phenxint.com

VS/C1-250-SSD

Phoenix International's Conduction/Convection Cooled VME Data Storage Blade delivers high capacity, high performance data storage for military, aerospace and industrial applications requiring rugged, extreme environmental and secure mass data storage.

This 6U single-slot module houses one or two each 6Gb Serial Attached SCSI (SAS) or Serial ATA (SATA) Solid State Disks or Rotating Hard Drives. These high speed modules will sustain write speeds in excess of 550MB/sec and 80,000 4KB random write IOPS. The VS/C1-250-SSD is available in an operating temperature range from -40 to +85 degrees C and functions at an altitude greater than 80,000 feet. It also complies with current defense department security standards, providing multiple levels of secure erase techniques.

The VS/C1-250-SSD's outstanding performance and versatility are enabled by Phoenix International's state-of-the-art   technology, which provides high transfer and I/O rates, endurance and data integrity.

FEATURES

› Serial ATA (SATA) and Serial Attached SCSI (SAS) Device Interfaces
› Operational Temperature Range From -40 to +85 Degrees C
› Up to 80,000 Feet Operational Altitude
› Sustained R/W Data Transfer Rates to 550Gb/sec
› 80,000 4KB Random Write IOPS
› Integrated SLC, eMLC or MLC NAND Flash Solid State Disk Drives
› Supports 7.2K, 10K or 15K RPM Hard Disk Drives
› Meets Military and IRIG 106-07 Declassification Standards
› Advanced Flash Management for Enhanced Reliability and Durability
› Low Power Consumption

For more information, contact: info@phenxint.com   **www.mil-embedded.com/p43730**

Mass Storage: Solid State Disk (SSD)

RunCore Co., Ltd.

6th Floor, No. 662 Lugu Avenue • Yuelu District, Changsha, Hunan Province P.R.C.
www.runcore.com/en/

RunCore Invincible PXIe Interface SSD

RunCore PXIe solid-state memory card not only has the advantages of high-speed read and write and large data capacity, but also has a vibration-resistant character at high and low temperatures for high reliability. Since it is an industry-standard PXIe transmission interface, it can be widely used in industrial control computers. Using a PCIe-SATA RAID chip and PXIe backplane for communications and data transmission, it is also used as an industrial-grade SSD controller to manage NAND flash devices. The on-board flash devices are divided into 32 channels, operating in parallel to deliver capacities up to 480GB. The entire board can be optimized to more than 450MB/s write and read speed, and it is widely used in a variety of high-performance storage.

For more detailed product information, please contact us at the following emails to discuss your application with a RunCore representative:

USA: sales@us.runcore.com
Europe: sales@eu.runcore.com
Asia-Pacific region: sales@asia.runcore.com

FEATURES

› Ultra High Performance
› Large Capacity
› Data Security
› Wide Temperature Shock

For more information, contact: sales@runcore.com   **www.mil-embedded.com/p53132**

Military Embedded Systems | Resource Guide

Phoenix International

812 W. Southern Ave. • Orange, CA 92865 USA
714-283-4800
www.phenxint.com

RPC12 Ruggedized Data Storage Array

The RPC12 is a high performance Fibre Channel/SAS/iSCSI host interface, 6Gb SAS/SATA-II RAID subsystem utilizing Hard Disk and Solid State Disk devices that deliver a level of operational environmental capability not previously available in COTS Data Storage Systems.

The Storage Area Network (SAN)-ready RPC12 features a 12-drive Array housed in a rugged 3U (5.25") panel height enclosure, providing 8Gb FC, 6Gb SAS or 10GigE iSCSI host interfaces to high performance 6Gb SAS and/or SATA SSDs and HDDs. It is designed to be compliant with military and industrial specifications such as MIL-STD-810F, MIL-STD-901D, MIL-STD-461F and NEBS Level 3.

The unique design of the RPC12's rugged, cableless, passive midplane-based, high density 3U chassis provides an increased environmental operational envelope (-20 to +70 degrees C, 45,000 ft. altitude with SSDs), redundant, hot-swap components and massive storage capacity, while assuring the highest level of data availability.

FEATURES

> Single Active or Dual Active, Failover/Failback Controller
> 8GB Fibre Channel, 6GB SAS or 10GigE iSCSI host ports
> Sustained Data Rates in Excess of 700MB/sec per port
> Battery-free Cache Backup
> Enclosed and Electronically Isolated Hot-swap Drive Canisters
> 40 to 440Hz, 90/240VAC Input Operation
> Linux, Windows and UNIX Support (VMware and Cluster Certified)
> Redundant, Hot-swap Components/FRUs
> Management GUI and Failover Software
> Very Cool Operation – Less than 12 Degress F Max. Temp. Rise
> Made in the USA

For more information, contact: info@phenxint.com

www.mil-embedded.com/p35575

Pinnacle Data Systems, Incorporated

6600 Port Road • Groveport, OH 43125 USA
Tel: +1 (614) 748-1150 • Fax: +1 (614) 748-1209
www.pinnacle.com

ComputeNode™ CompactPCI Chassis Products

PDSi's ComputeNode line offers a range of NEBS Level 3-compliant CompactPCI chassis in sizes from 1U to 4U. These carrier-grade chassis include a horizontal design, superior air cooling, cPCI and cPSB (PICMG 2.16) backplanes, redundant hot-swappable fans, hot-swappable, front-accessible AC or DC power supplies, and rear single or dual power feeds. All 2U and larger ComputeNode platforms include PDSi's unique Alert!Node™ (or Enhanced Alert!Node) alarm card, an intelligent chassis management controller for comprehensive fan and power monitoring. The Alert!Node card does not occupy a CompactPCI slot, front or rear.

OEMs and Independent Software Vendors (ISVs) can also take advantage of PDSi's design, integration, and support services, including custom board and system design, validation and certification, production assembly and test, as well as extended service programs.

FEATURES

> Proven NEBS Level 3-compliant design
> Designed for high-availability applications
> Redundant hot-swap fans
> Redundant AC or DC power supplies
> Power filters and dual-feed power
> Easily serviced
> Broad range of chassis choices: 1, 2, 3, and 4U sizes
> cPCI and PICMG 2.16 (cPSB) backplanes
> Alert!Node intelligent chassis manager
> Customization and integration services available

For more information, contact: info.sales@pinnacle.com

www.mil-embedded.com/p44295

SIE Computing Solutions, Inc.

10 Mupac Drive • Brockton, MA 02301 USA
508-588-6110
www.sie-cs.com

COMPUTING SOLUTIONS

716 Series Conduction Cooled ATR Enclosures

716 Series offers a wide range of COTS solutions from a rugged, precision-machined design. Engineered for strength, light weight, and maximum cooling in a conduction-cooled environment, the 716 Series incorporates a unique frame and configurable conducting walls that allow the ATR to be tailored to meet a wide range of thermal requirements.

FEATURES

› Precision-machined construction
› Available in 3U or 6U card formats
› Rugged deployment
› Expansive range of ARINC sizes
› Modular power supply
› AC or DC filtered inputs
› System performance monitoring
› Multiple bus architectures
› Cold-start heaters
› Configurable I/O panel

For more information, contact: info@sie-cs.com   **www.mil-embedded.com/p34423**

Elma Bustronic

44350 Grimmer Blvd. • Fremont, CA 94538 USA
510.490.7388
www.elmabustronic.com

bustronic
Elma Bustronic

OpenVPX Backplanes

Elma Bustronic is the leader in OpenVPX backplanes with unsurpassed knowledge, leadership, and innovation in the industry.

Working with top VPX industry experts, Elma Bustronic has developed solutions geared for intuitive module configurations. Profile options include 3U, 6U, and hybrid versions. Our signal integrity studies will help you improve performance for your custom design. Have a challenging design or need someone to guide you through the specification? Come to the leader, Elma Bustronic.

Elma Bustronic's VITA product line also includes a wide range of VXS backplanes in Star, Dual Star, and Switchless Mesh configurations, as well as standard and customized VME/VME64 backplanes and accessories.

Contact us today for more information on these products or our customization capabilities.
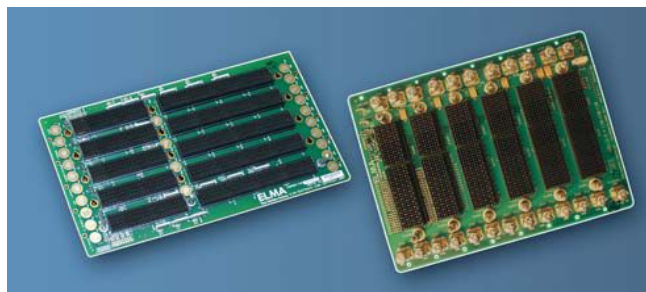
FEATURES

› Leader in OpenVPX backplane profiles
› OpenVPX backplanes to 6.25 Gbaud/s data rates
› Wide range of sizes and configurations in 3U, 6U, hybrid, development, and power and ground
› Large selection of accessories available including load boards, SerDes test devices, RTMs, and extender boards
› High reliability, consistency, and superior manufacturing and testing capabilities
› Signal integrity analysis experts for high-speed designs
› VPX cabling options available
› Custom solutions available

For more information, contact: sales@elmabustronic.com   **www.mil-embedded.com/p45467**

Vector Electronics & Technology, Inc.
11115 Vanowen Street • North Hollywood, CA 91605 USA
800-423-5659
www.vectorelect.com

VME/VME64x Backplanes and Series 790 Chassis

VME/VME64x Backplanes per ANSI/VITA 1.1-1997

MIL-1-46058C conformal coating available for all Vector backplanes

Vector VME64x 6U monolithic backplanes are a "true" 6U footprint, fully RoHS-compliant and available for immediate shipment from our factory.

Vector VME64x monolithic backplanes for 6U cards are available in 2, 3, 4, 5, 7, 8, 10, 12 and 21-slot configurations. 'True' 6U (10.317" H) for easy replacement or addition to any 6U rack system.

Custom backplane design services available.
Please contact
inquire@vectorelect.com.

Series 790 Chassis

The Series 790 is a ruggedized version of the Series 730. It has been designed and tested to MIL-STD 461D for harsh environmental conditions. RFI/EMI frequency protection at high and very low bandwidth interference, condensing humidity, and rigorous shake and drop tests are the strong points of this design. A removable rear panel section allows custom I/O connector panelization at minimum cost. This series is also available at a much lower cost than competitive units with the same features.

Go to www.vectorelect.com
to order our New 2011 Catalog.
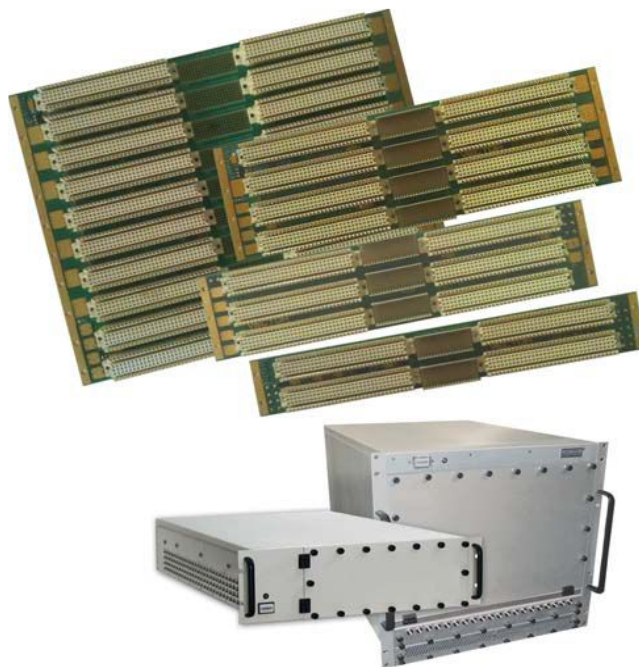
FEATURES

› **VME/VME64x Backplanes**

– Greater than 64MHz high-speed design
– 10-layer FR-4 construction, UL 94-VO, RoHS-compliant
– Signal lines shielded, controlled impedance
– OR-logic electronic daisy-chaining (EBG) with on-board termination
– SMT capacitors, resistors, etc. replace through-hole for increased reliability
– Screw terminal power connections

› **Series 790 Chassis**

– Well suited for Military / Aeronautics
– VME64x, VME64 or cPCI-compatible
– MIL-STD-461D, E for shock/vibration, EMI/RFI
– CE and FCC Class B, Part 15
– Many power options

SIE Computing Solutions, Inc.

10 Mupac Drive • Brockton, MA 02301 USA

508-588-6110

www.sie-cs.com

VPX/OpenVPX Series VITA 46/48/65 Backplanes

SIE Computing Solutions' VPX backplanes are designed to the latest VITA 46, 48 and 65 standards. SIE OpenVPX backplanes utilize the BKP3-CEN06-15.2.2-3 profile ratified by VITA 65, and are also available in custom configurations. The 5-slot full mesh 3U VPX and 6-slot 3U OpenVPX backplanes are designed for a wide array of VPX/OpenVPX applications. The highly configurable SIE backplane line offers maximum bandwidth in a compact size. The 3U VPX backplane provides greater I/O flexibility through I/O PLUS™, an innovative use of configurable I/O daughtercards that accommodates an array of VPX applications.
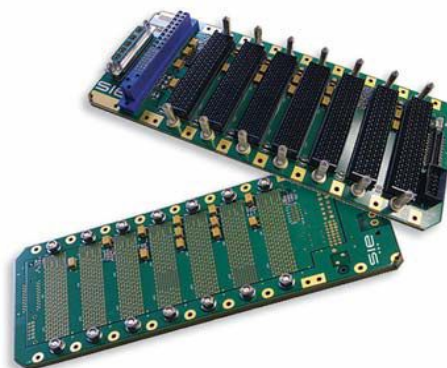
FEATURES

› Designed for board-agnostic system design and integration

› Range of full mesh VPX and OpenVPX configurations and profiles

› 2 slots dedicated I/O daughtercards

› VITA 46, 48 and 65-compliant

› Over 200 watts per slot

› 28-layer board

› RoHS-compliant product features

For more information, contact: info@sie-cs.com

www.mil-embedded.com/p45468

Packaging/Mechanical Chassis: Rugged Chassis

LCR Electonics, Inc.

9 South Forest Ave • Norristown, PA 19401 USA

800-527-4362

www.lcr-inc.com/systems-backplanes/index.html

Six Pack Rugged Military ATCA System

The Six Pack ATCA System is a rugged 6U rack-mountable chassis designed for military use. The 6-slot mesh ATCA backplane allows for two Vadatech shelf managers. The six fans in the fan tray (189 CFM per fan) provide redundant cooling up to 300 W per slot and provide active RTM cooling.

The 0.075 steel outer shell, machined card guides and other rugged features are designed to withstand harsh environments, including shock and vibration, while operating over a temperature range from 0°C to 70°C. The Six Pack ATCA System can be provided with redundant AC power supplies or DC power entry modules.

The Six Pack is part of LCR's family of more compact ATCA products. This chassis family, along with the larger 14-slot family, has been incorporated in platforms that cover all aspects of mission-critical computing, such as weapons control; command and control; and intelligence, surveillance and reconnaissance; and it is perfect for airborne, ground mobile, naval and UAV applications.

FEATURES

› High-performance for networking and communication requirements

› Capabilities and processing power comparable to rack-mount servers at one half the weight and improved field mobility

› Bladed ATCA architecture incorporates multiple modules into a single chassis with a common enclosure; fans and power supplies weigh significantly less than rack-mount server applications

› ATCA rugged off-the-shelf system available in 4U to 6U heights with 3- to-6 slot full mesh ATCA backplanes

› Redundant AC 1500 W power supplies or -48 V DC redundant rear power entry

› Meets MIL-STD-810 temperature, shock, vibration and acceleration standards

For more information, contact: military-sales@lcr-inc.com

www.mil-embedded.com/p53140

Military Embedded Systems   Resource Guide

Parvus Corporation

3222 Washington Street • Salt Lake City, UT 84115 USA
800-483-1533
www.parvus.com

ACS-5180 – PC104+ 80 Watt – MIL-704/1275 Power Supply

The ACS-5180 is a rugged PC/104-*Plus* isolated power supply capable of supplying up to 80 watts of power output over extended temperature ranges to embedded computing devices in demanding military/civil ground vehicle, shipboard, and aircraft applications. Designed for extended temperature operation (-40 to +85C per MIL-STD-810G) and demanding power conditions dealt with onboard military ground vehicles (MIL-STD-1275D) and aircraft (MIL-STD-704F) installations, this highly efficient (90%+), filtered DC/DC converter supplies DC voltage outputs (+3.3V, +5V, +12V) over the PC/104 (ISA) bus, PC/104-*Plus* (PCI) bus, or onboard screw clamp terminal.

Featuring a rugged mechanical design, this small form factor (3.550" x 3.775") card is designed to be used as the bottom card in a PC/104 system stack, operate without heatsinking or any active cooling, and provide resistance to high levels of shock and vibration. Featuring an onboard MIL-STD-461 EMI filter and transient protection for the 250V spike and 100V surge requirements of MIL-STD-1275, the card can replace the typical in-line power conditioning/EMI filtering devices integrated into military embedded systems.
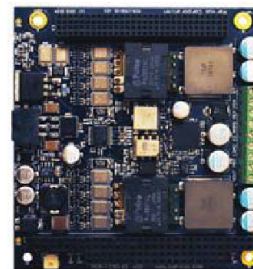
FEATURES

› PC/104-*Plus* Power Supply
› 24/28V DC-DC conversion for PC/104, PC/104-*Plus*, PCI-104, and other small form factor apps
› Compliance with MIL-STD-1275, MIL-STD-704, MIL-STD-461
› Designed to Meet MIL-STD-810G, MIL-STD-1275D, MIL-STD-704F, and MIL-STD-461F
› 1500V Galvanic Isolation DC
› Up to 80 watts at +85C: +5V @ 16A; +12V @ 2.0A, +3.3V @ 8.0A
› Reverse Polarity, Voltage Transient, Surge, Spike, Over Current
› PC/104 Bus (ISA), PC/104-*Plus* Bus (PCI), Screw Clamp Terminal
› Remote Shutdown Support

For more information, contact: sales@parvus.com   www.mil-embedded.com/p45654

Excalibur Systems

311 Meacham Ave. • Elmont, NY 11003 USA
Tel 800-MIL-1553 • Fax 516-327-4645
www.mil-1553.com

MACC – Miniature Airborne Communications Converter

Excalibur's MACC is a complete Miniature Airborne Communications Converter designed to convert bidirectional messages. Small, light, and with low power dissipation, the MACC is the perfect solution for airborne communications incompatibilities.

When designing or upgrading an avionics system, engineers are often faced with integrating components that were not originally designed to work together. For example, a flight computer that communicates using RS-232 may need to control a component that expects MIL-STD-1553 messages, or an INS may transmit ARINC 429 labels needed by a multifunctioning display designed around RS-422. No longer will you have to choose between giving up one of the units because of communications incompatibility, or adding another computer system to convert between the two. Excalibur's MACC converts messages between:

- MIL-STD-1553
- ARINC 429
- Serial (232/422/485)
- Ethernet
- Discrete

FEATURES

› Conduction cooled
› Waterproof aluminum case
› Black, anodized casing with built-in heat sink
› No moving parts – fanless
› Industrial temperature range
› Uses MIL-38999 I/O connectors
› Rugged, screw-down internal assembly for shock

For more information, contact: excalibur@mil-1553.com   **www.mil-embedded.com/p53031**

TE Connectivity

PO Box 3608, MS 38-55 • Harrisburg, PA, 17105 USA
Toll Free: 1-800-522-6752
www.te.com/ADM

FORTIS ZD

TE's advanced FORTIS ZD high-speed backplane connector is designed for military and commercial aerospace applications. As demands on systems for real-time intelligence intensify, the importance of high-performance interconnection becomes critical. TE set out to design a new connector system that supports increasing bandwidth requirements in a ruggedized format to withstand the increased shock and vibration requirements of emerging military applications.

FEATURES

› Allows 10 Gb/s+ data rates
› Extreme mechanical and electrical performance for the most demanding applications
› Modular design allows for user configurability and modular evolution
› Three shell varieties for application versatility
› Mini-Box contact system provides 4 points of contact for ultra reliability
› www.TheFutureUnleashed.com for more information

For more information, contact: product.info@te.com   **www.mil-embedded.com/p52822**

Rugged Computer Systems: Electronics Packaging

TE Connectivity
PO Box 3608, MS 38-55 • Harrisburg, PA 17105 USA
Toll Free: 1-800-522-6752
www.te.com/ADM

Mezalok

TE Connectivity's Mezalok connector is a high-reliability mezzanine connector that more than doubles the speed and durability of competing technology, making it the most viable option for today's military and commercial aerospace applications.

The need for embedded electronic systems in rugged platforms has been continually increasing and driving more stringent requirements for board-to-board connectivity. TE engineers responded to these needs by developing a super-fast, super-durable and super-redundant mezzanine connector ideal for extreme environments.
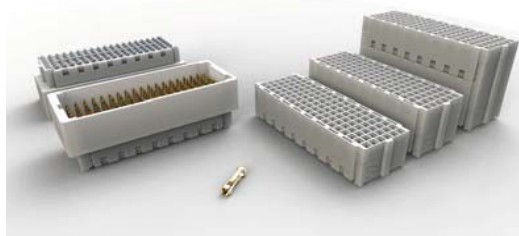
FEATURES

> 60 and 114 positions

> Mini-Box contact system provides 4 points of contact for ultra reliability

> LCP plastic housings offer superior thermal stability and are low-outgassing

> 114 position footprint compatible to XMC footprint

> 500 mating cycles durability

> www.ElementalAdvantage.com for more information

For more information, contact: product.info@te.com

www.mil-embedded.com/p46728

Rugged Computer Systems: Flat Panel System/Combat Display

CRI, a Division of RSI, Inc.
1670 Kohler's Crossing • Kyle, TX 78640 USA
512-268-7500 • 1-806-872-2700
www.cri-rugged.com • www.rsi-cri.com

ISO 9001:2008
AS 9100
AS 9120
BSI Registered
FM 45531

Battlefield Systems

CRI has provided cost-effective, high-quality, custom rugged COTS solutions for military, aerospace, and industrial applications since 1994. CRI equipment is designed to MIL-STD-810, MIL-STD-461, MIL-STD-167, and MIL-STD-704 and has a proven record in harsh environments. CRI's varied product line of rugged equipment includes:

• **Displays** (AC/DC, Rack, Drawer, touch-screen, NVIS, Sunlight Readable, Low Temp.)

• **Servers** (AC/DC, Rack, stand-alone, Size: 1-5U)

• **KVM/Ethernet Switches**

• **Routers**

• **Drive Storage Enclosures**

• **Power Supplies** (AC/DC 400Hz)

• **Panel PCs**

RSI est. 1983 • CRI est. 1994

FEATURES
> **Rucksack – Portable Workstation**
  – Multi-use chassis
  – Integrated 6U cPCI or VME backplane
  – Integrated slot-load DVD
  – Cross-ventilated chassis for maximum thermal performance
  – Front panel indicators for internal KVM and power available
  – Front panel LCD controls
  – Keyboard panel with trackpad
  – Custom colors or camouflage
  – MIL spec power connectors
  – 350W power supply, AC/DC input
  – Replaceable fan filters

For more information, contact: info@rsi-cri.com

www.mil-embedded.com/p53135

IBI Systems, Inc.

6842 NW 20th Avenue • Fort Lauderdale, FL 33309 USA
954-978-9225
www.ibi-systems.com

ST-9020

ST-9020 is a Rugged Computer System with built-in 20.1" high-resolution display.

It utilizes a passive backplane and plug-in CPU card for ease of upgradability and serviceability. The plug-in cards are secured with card clamps. All the drives are shock isolated. The hard drives can be fixed or removable. The system can be configured as per your requirements.

Please contact us for all your Rugged Computer System requirements.

FEATURES

> Rugged 19" rack-mount chassis
> 14-slot ISA/PCI passive backplane
> 460 Watt 110/220V AC 50-400Hz power supply
> Pentium 4, 3.4GHz CPU card with 2GB ram
> 20.1" color TFT display with 1600 x 1200 resolution and true colors
> Shock-isolated DVD/CD-RW drive
> 73GB Ultra SCSI removable and shock-isolated drive
> Size: 15.72"H x 19"W x 19"D
> Meets MIL-STD-167-1 vibration
> Meets MIL-STD-901 light weight, Grade A, Class 1 shock
> Windows XP Pro or Windows 7

For more information, contact: ibisys@aol.com   **www.mil-embedded.com/p53152**

IEE, Inc.

7723 Kester Ave. • Van Nuys, CA 91405 USA
818-787-0311
www.ieeinc.com/military-products

Advanced LCD Military Displays

IEE has an extensive product portfolio of MIL-STD-qualified flat panel displays, keypads and switch bezels that feature lightweight enclosures, video processing, serial communication and power conversion for rugged environments. Our proven expertise in integrated display design includes features such as multi-mode backlighting, brightness control from NVIS to high bright, optical bonding, touch screens, EMI filtering, custom anti-reflective/anti-glare (AR/AG) coatings and rugged packaging.

Recent products include low-power integrated ARM processors that support embedded Linux and Android operating systems with a wide range of communication peripheral interfaces. All of IEE's products are manufactured in our 65,000 sq. ft. modern manufacturing facility in Van Nuys, Ca. With rapid prototyping through high-volume production services available, we support our customers' efforts to reduce cost and lead-times.

FEATURES

> Enhanced displays from 3.5" to 24" and beyond; standard sizes include: 3.5", 7", 10.4", 12.1", 15", 19" and 24"
> Touch-screen integration: resistive: 5-wire and 8-wire, multi-touch; IR, SAW, and projected capacitance
> Active brightness enhancement: high bright LED backlights/controllers
> Passive brightness enhancement films: BEF, DBEF, ESR
> Automatic brightness control and on-screen display controllers
> LCD video controllers: DVI and analog RGB inputs are standard, with optional support for composite video, Ethernet and USB
> Products meet or exceed military standards for environmental, EMI/EMC, construction, lighting and power

For more information, contact: sales@ieeinc.com   **www.mil-embedded.com/p53145**

IEE, Inc.

7723 Kester Ave. • Van Nuys, CA 91405 USA
818-787-0311
www.ieeinc.com/military-products

Hand Held Control Display Units

IEE's family of Hand Held Control Display Units (HH-CDUs) are combat-proven, sunlight-readable, MIL-qualified units designed for soldier-mount, vehicle-mount or manpack handheld use. The backlit enhanced displays are available in sizes and resolutions ranging from 4-line, 2.8" character displays to full graphic displays from 3.5" QVGA through 7" WSVGA.

The product family employs a wide range of low-power, high-performance reprogrammable microprocessors supporting industry standard peripherals and communication protocols. The device operator interfaces are matched to the application-specific I/O needs through buttons, keypads, and/or touch-screen technologies responding to stylus, finger, or gloved-hand inputs.

The HH-CDUs provide functionality ranging from basic text I/O to full web-browser features supporting HTML 5.0 executing Java scripts. Our products include the Linux or Android OS hosted on an ARM processor and execute off-the-shelf open-source applications.

FEATURES

› Military-qualified, combat-proven, rugged units for MIL-STD-810 environments, including wide -40°C to +70°C operating temperature
› Display sizes and resolutions range from 4-line 2.8" character displays to full graphic displays from 3.5" QVGA through 7"-wide SVGA
› Multi-mode backlighting systems that provide sunlight readability and NVIS compatibility
› Integrated touch-screen technology: resistive, IR, SAW, multi-touch
› Low-power, high-performance reprogrammable microprocessor-based operation supporting customization
› Flexible serial data interfacing through industry standard Ethernet and EIA RS-422
› Optional "zeroize" switch under an accident-proof cover

For more information, contact: sales@ieeinc.com   **www.mil-embedded.com/p53144**

Emphase

35 Thompson Street • South Burlington, VT 05403 USA
802-735-1799
www.emphase.com

2.5" SATA SSD S6 – MIL-SPEC-Grade

The Emphase MIL-SPEC S6 Series SSD is engineered to withstand the severe conditions encountered in the line of duty while delivering comprehensive data protection. An SLC flash device available in a 2.5" footprint and SATA (3 Gb/sec.) bus interface, and the MIL-SPEC S6 SSD has one clear objective: Secure data. This vault-like device can process data with read speeds up to 220 MB/sec. and write speeds up to 180 MB/sec. Mission-critical applications have no room for error, which is why the MIL-SPEC S6 SSD is designed to meet the environmental requirements of MIL-STD-810F.

The Emphase MIL-SPEC S6 SSD expertly manages data in breach of security scenarios, including data write protection, secure and quick erasure, and data destruction. Data can be eliminated in less than two seconds through an ATA command or with the push of a button. Should the drive lose power during a protect, erase, or destroy command, the device will resume the command as soon as power is restored. Other features include internal DRAM cache, SuperCap power protection, and custom data encryption.

FEATURES

› Thermal Shock: MIL-STD-810F 503.4 Procedure II (-40°C~85°C, 10 cycles)
› Drop: MIL-STD-810F 516.5 Procedure IV (Transit Drop Specification)
› Shock: 1500 G/MIL-STD-810F 516.5/40 G, 11 ms, 18 shocks (3 shocks for ± X, Y, Z axis)
› Vibration: 20 G/MIL-STD-810F 514.5C-1
› Altitude: 80,000 ft.
› Operating Temp: Standard, 0°C~70°C; Wide Temperature -40°C~85°C
› Erase and Destroy: USA-AF AFSSI 5020; DoD 5220.22-M; USA Navy NAVSO P-5239-26; NSA 130-2; USA Army 380-19; NISPOMSUP Chap 8, Sect. 8-501
› Applications: Military/Defense Server, Tactical, and Data Systems

For more information, contact: info@emphase.com   **www.mil-embedded.com/p53137**

Rugged Computer Systems: Mission Computer

Black Diamond Advanced Technology

7450 South Priest Drive • Tempe, Arizona 85283 USA
Phone: 855-855-BDAT • Fax: 480-705-4216
www.bdatech.com

Modular Tactical System

The Modular Tactical System (MTS) centralizes communications, equipment and power management into a wearable C4ISR OTM (Command, Control, Computers, Communications, Intelligence, Surveillance, and Reconnaissance On The Move) and precision targeting tool that delivers true foot-mobile computing. With a design that anticipates how an operator moves, strategizes and acts, the MTS funnels control of mission-critical peripherals to a central display to eliminate battery and equipment redundancies and lighten load-out.

In the vest-based configuration, a small, low-profile Tactical Mission Controller – containing processing, peripheral control and power management functions – is carried on the back. A flip-down pocket on the front of the vest allows quick access to a 6.5-inch sunlight- and NVG-viewable display. All cables are routed through an interchangeable cummerbund that integrates body armor, is secured to plate-carriers, and provides non-obtrusive access to I/O ports. The MTS can also be employed in bag or vehicle configurations.

BLACK DIAMOND
ADVANCED TECHNOLOGY

FEATURES

> Transition in seconds from system operation to direct combat
> Operate while walking
> Control fielded equipment from a single display
> Run all gear from one battery
> Reduce weight carried and eliminate redundant equipment by integrating radios and VDL devices
> IP67 – Operate in a dust storm or up to 1 meter of water
> Network 2 or more radios simultaneously
> **Touch screen display can be operated with gloves**
> **Ensure unrestricted access to ammo**

For more information, contact: info@bdatech.com   www.mil-embedded.com/p52871

Rugged Computer Systems: Mission Computer

CES – Creative Electronic Systems

Avenue Eugene-Lance 38 • 1212 Grand Lancy, Geneva Switzerland
+41.22.884.51.00
www.ces.ch

Mission Computer

Mission computers for UAV systems provide high-level flight management, mission management and payload management functionalities. Processing power is provided by one or more CES processor boards, which communicate over VME/VPX and/or high-speed serial links. Custom processing or I/O functionality can be implemented in FPGAs.

A large variety of avionic interfaces is available from CES, including MIL-1553, ARINC 429, RS-422/485 and many others. An integrated Gigabit Ethernet switch assures standard connectivity, both internally and externally, in addition to the VME or VPX backplane. Multiple high-speed serial links can be used for additional high-bandwidth connectivity. Video processing capabilities include analog and digital video input and output, still image and streaming video compression, as well as custom processing in the processor or FPGA. Flash-based mass storage is available for mission data storage or video recording.

CES
CREATIVE ELECTRONIC SYSTEMS

FEATURES

> Safety-critical and certifiable solution
> COTS building blocks
> Powerful and rugged multiprocessing architecture
> Custom processing in the FPGA
> Wide range of avionic interfaces
> Multiple Gigabit Ethernet interfaces
> Video input, output and compression
> Flash-based mass storage
> Sealed ATR military enclosures
> DAL A – DAL C certified

For more information, contact: ces@ces.ch   www.mil-embedded.com/p30050

CePOINT Networks, LLC
1W Otterson Street • Nashua, NH 03060 USA
Tel: 603-883-7979 • Fax: 603-883-3266
www.cepoint.com

Manufacturers of Rugged Portable & Airborne DVR systems w/IRIG-B time stamp

Studio9000™ DVR IRIG-B

Real-time Digital Video Recorder (DVR) system for robust scientific image acquisition and analysis

Studio9000 DVR system performs with blazing speed, featuring uncompressed (or compressed) real-time video capture and recording with optional precision IRIG-B time stamping and GPS interface capabilities. Standard digital or composite analog video acquisition in color NTSC/PAL, SECAM, RGB YCrCb 4:2:2, or in monochrome format – CCIR (625 lines) and EIA (525 lines) – are supported. Optional SDI is also supported. Up to 240 fps (analog), and very high-speed digital video up to 1280 x 1024 resolution, and 30 fps up to 500-1,000 fps (digital) is possible. Other features include: simultaneous capture/playback of four video streams; up to two or more channels of real-time simultaneous record and play; unlimited multicam editing and reediting of captured video without degradation or frame loss; captures continuous real-time video directly to system hard disk or memory; compact, rugged 2RU, 3RU, or 4RU MIL-COTS format; capture and stream directly to disk at up to 528 MBps; capture directly to system hard drive from different video formats and sources supported by Studio9000 DVR; monochrome or color at 8 bits, 10 bits, 12 bits, 14 bits, and more, including area scan, progressive scan, and line scan. Optional interface features include analog BNC, digital LVDS, Camera Link, USB, and 1394 FireWire cameras.

Applications:

- Airborne video recording
- Object tracking and time reference measurement
- Missile range testing
- Endless video program looping
- Security recorder/player
- Bullet explosion testing
- Industrial monitoring
- Portable field production
- Desktop video capture station
- Surveillance recorder

Studio9000 DVR greatly simplifies the process of time referencing object position and timing measurements by integrating real-time video acquisition, real-time IRIG time stamp, and GPS position data.

FEATURES

> Capture continuous real-time video directly to hard disk at up to 528 MBps; 8-bit, 10-bit, 12-bit, 14-bit, 24-bit mono or color

> Analog RS-170, NTSC/PAL, RGB, and digital LVDS, Camera Link, USB, FireWire 1394, and RS-644 or RS-422 camera interface options

> Video resolution: 640 x 480, up to 1280 x 1024 pixels; compressed or uncompressed video formats include: AVI, MJPEG, optional MPEG-4

> Digital clock circuitry; capture high-speed, high-resolution images from RGB or composite; progressive scan, line scan, and area scan

> Optional SDI video I/O (SMPTE 259M, 270 Mbps) with embedded AES/EBU audio

> IRIG-B and GPS formats include: Time code generator, IRIG receiver, ANT BNC input connector, and DB-9-pin RS-232 connector

> Real-time simultaneous capture of up to four channels; stream video directly to hard drive, memory, or display output

> RAID 0 storage with capacity up to 4.8 TB optional, and expandable with CePOINT's optional NAS RAID storage for extended duration of video

> External event triggers; up to 4- or 8-channel digital I/O for programmable triggers

> External interface ports include: RJ45 Ethernet, 1 x PS/2 keyboard, 1 x PS/2 mouse, VGA, RS-232, or RS-422

> Support for Region of Interest (ROI) video manipulation, packed and planar; YUV 4:2:2

> Rugged MIL-COTS format; lightweight, rugged 19" 2U, 3U, or 4U rack-mount, airborne or portable with 24 V or 28 VDC option

Crystal Group Inc.
850 Kacena Road • Hiawatha, IA 52233 USA
319-378-1636
www.crystalrugged.com

SS16 – Fully sealed rugged server

The Crystal SS16 sealed server offers all aluminum construction with custom EMI gasketing to seal the server chassis. I/O is facilitated through quick-locking mil circular connectors. The SS16 sealed server offers immersion protection rated to IP67. Convection cooling enables an extended temperature range of -30C to +55C, with a -40C SSHDD option. Cold-plate mounting is available for higher temperature environments. The SS16 sealed server is tested to MIL-STD-810F and designed to MIL-STD-461F.

In environments where performance, ruggedness and reliability are imperative, the SS16 is an ideal fit. This rugged, fully sealed server was designed for embedded shipboard, airborne, and land-based applications. As always, each Crystal product comes with a standard three year warranty. Extended warranties and configuration management are also available upon request.

The SS16 utilizes a Mini-ITX motherboard with Intel's i7/i5 processors and a QM57 chipset. The design is easy to upgrade to the next innovative Mini-ITX design, allowing customers to stay on the forefront of computational capabilities years ahead of other platforms.

The SS16 provides a leading edge COTS-based open architecture compute platform that is beyond rugged; it approaches bulletproof. Please download our test results at www.crystalrugged.com/products/embedded/rugged-sealed-server-ss16.aspx.

_____

Crystal Group Inc., an employee-owned small business located in Hiawatha, Iowa, USA, designs and manufactures deployment-ready, rugged servers, displays, networking devices, embedded systems, and storage devices that fit mission-critical applications. For over 24 years, Crystal Group has been developing rugged solutions for our customers' most challenging needs. Crystal excels at providing commercial-off-the-shelf (COTS) and custom-designed computer systems tailored to specification. All products are designed and built to meet and exceed MIL-STDs 810F/G, 167-1, 461E/F and MIL-S-901D. Crystal Group also offers integration services in addition to configuration management, product life-cycle planning and three-plus year warranties. Contact a Crystal Business Development Manager for more information: 1.800.378.1636.

The Most Trusted Provider of Computer Architecture for
Operational, Deployable,
and High-Reliability Applications

FEATURES

› Rugged sealed server rated to IP67
› Up to six (6) 2.5" SSD hard drives
› Intel Core i7/i5/Celeron micro-FCPGA 989 Mini-ITX with QM57 chipset, VGA, four COM, Dual LAN, one PCIe x16, 8GB DDR3
› Weight 23 lbs. max
› Extended temperature range, -40C to +55C, natural convection cooled, extreme temperature range
› Designed to MIL-STD-810F temp., altitude; MIL-STD-1275D; MIL-STD-704E, MIL-STD-461; MIL-STD-810F vibe, shock
› Configuration management available; designed for long-life through Intel® Roadmap partnership
› Mounts via top or bottom surfaces
› 6061-T651 strain-hardened structural aircraft aluminum ensures battle endurance
› MIL-C-5541E and powder coat finish provides extended corrosion resistant life
› Hermetic mil circular connectors ensure connections in harsh environments
› Multiple attachment points on motherboard and CPU heatsink stabilize PWB preventing BGA damage
› All fasteners and connectors retained with locking mechanisms to withstand extreme shock and vibration
› Natural convection chassis with extended fin surfaces allows extreme temperature operation
› Internal conduction and external convection cooling techniques extend COTS electronic performance
› Completely sealed for radiated susceptibility and radiated emissions; filtered for conducted susceptibility and conducted emissions
› Finish applied maintains superior EMI/EMC characteristics throughout life of the unit
› Superior grounding paths including external ground stud for noise-free operation

Rugged Computer Systems: Mission Computer

Kontron

14118 Stowe Drive • Poway, CA 92064-7147 USA
858-677-0877
www.kontron.com

Rugged Systems & Enclosures

Kontron provides a unique blend of both COTS and specialized systems designs that can be custom-tailored to meet the most demanding requirements of our customers and program partners. All of Kontron's advanced chassis designs are optimized for both flexible configurability and rugged performance, with a range of choices in each design: from cooling methods to bus/backplane architectures to processors, I/O and power supplies.

In addition, Kontron has a wide range of rugged enclosures for a variety of applications – providing standard, semi-custom as well as full-custom enclosures. Kontron's knowledge and experience goes beyond MIL-SPEC and DO-160 requirements and design with strength, serviceability, and modularity in mind, offering the newest technologies to our customers, who deploy systems on a wide range of military platforms. Rack-mount or ATR form factor, shock isolated or hard-mount, Kontron enclosures are tuned to a variety of shock, vibration and EMI standards based on operating environment.

FEATURES

› Expertise in ruggedization, certification, full system integration and total custom solutions

› Extended product life-cycle management with long-term support

› Largest selection of product hardware available, including rugged COTS boards and platforms for interoperable, standards-based solutions

› Strategic partnerships that leverage the latest technology immediately as it's available (Intel, Freescale, Wind River, etc.)

› Dedicated military teams located strategically around the globe

For more information, contact: info@us.kontron.com

www.mil-embedded.com/p53147

Rugged Computer Systems: Mission Computer

Parvus Corporation

3222 Washington Street • Salt Lake City, UT 84115 USA
800-483-1533
www.parvus.com

DuraCOR 830 – Small Form Factor Tactical Mission Computer

The DuraCOR® 830 is a rugged tactical mission processor based on the ultra low-power 1.6 GHz Intel Atom CPU, optimally designed for Size, Weight, and Power (SWaP)-constrained military/aerospace vehicle and aircraft installations. Mechanically compact (<4" tall, <4lbs.), this small form factor computer subsystem integrates Gigabit Ethernet and dual MIL-STD-1553 (optional) connectivity, along with digital I/O and standard PC interfaces, including analog video, audio, 3x USB, 2x RS-232 serial, keyboard, and mouse. For application flexibility, the unit comes with a removable 16 GB CompactFlash disk pre-loaded with a Linux or Windows Embedded operating system and an 8 GB industrial solid state disk (SSD) for additional file storage.

Designed for avionics and ground vehicle platforms to enhance situational awareness and computational capabilities, the DuraCOR 830 delivers processing and multimedia performance similar to Pentium M-based DuraCOR solutions, but with less power consumption, double the RAM memory capacity, and dual Flash SSDs.

FEATURES

› Intel Atom Z530 Processor, 1.6 GHz Clock Speed

› CompactFlash Solid State Disk 16 GB, Removable and Solid State Disk, 8 GB, Fixed

› 1x Gigabit Ethernet 10/100/1000 Mbps

› Up to 2x Open PC/104-*Plus* Slots for Expansion Cards

› Approx. 3.9" x 4.6" x 8.7" in Size

› MIL-STD-704E, MIL-STD-1275D Compliant

› MIL-STD-810G: Shock, Vibration, Temperature, Altitude, Humidity (Designed to Meet)

› MIL-STD-461E: Radiated and Conducted Emissions and Susceptibility (Designed to Meet)

For more information, contact: sales@parvus.com

www.mil-embedded.com/p47857

Octagon Systems

7403 Church Ranch Boulevard • Westminster, CO 80021 USA
303-430-1500
www.octagonsystems.com/products/rmb-c2.aspx

Rugged Mobile Computers

Octagon Systems has leveraged its 30 years of experience developing rugged, reliable industrial computers. Octagon Systems' powerful line of computing systems for mobile applications integrates wireless, LAN, serial and USB technologies in our rugged, fanless enclosure. These solutions let you focus on developing your application software without the distraction of having to find compatible components. Application software can be reused, easily porting from one system to another.

Modular design enables custom functionality with COTS convenience and without large upfront costs. Software drivers for Linux and Windows XP are options to speed your software development. We strive to provide each system with careful circuit design, component selection, testing and superior heat dissipation to maximize reliability and usability for your application. ISO 9001 certification guarantees a high level of quality in our products and service.

RMB-C2

The RMB-C2, currently Octagon's most rugged mobile computer, excels in harsh environments and applications where the cost of failure is extremely high. The field-proven RMB-C2 incorporates the best-in-class Octagon Hedgehog mobile power supply, a sophisticated shock and vibe system and a high density IP65 environmentally sealed extrusion to prevent dust and moisture entry.

The modular design enables custom functionaltiy with COTS convenience without large upfront costs. Software drivers for Linux and Windows XP are options to speed your software development. We strive to provide each system with careful circuit design, component selection, testing and superior heat dissipaton to maximize reliability and usability for your application. ISO 9001 certification guarantees a high level of quality in our products and service.

Designed, manufactured and certified in the U.S.

TECHNICAL SPECS

- Two USB 2.0
- One 10/100 BASE-T Ethernet
- CAN Bus
- Three RS-232, One RS-232/485 serial ports
- PC/104 and Mini PCI expansion

APPLICATION AREAS

- Military, Mining, Industrial

FEATURES

› Modular design ensures rapid customization

› Military connectors

› Designed, manufactured and certified in the U.S.

› Wide temperature range -40° to 75° C

› Internal power supply immune to extreme transients

› Application software can be reused, easily porting from one system to another

Visit our website
www.octagonsystems.com
to see our full line of mobile computers

FLEET

FLEET-Atom

FLEET-M

RadiSys Corporation

5445 N.E. Dawson Creek Drive • Hillsboro, OR 97124 USA
503-615-1100 • Fax 503-615-1115 • Toll-Free 800-950-0044
www.radisys.com

The RadiSys C2 Server

High Performance, Less Weight and Ruggedized

The C2 Server provides flexible high-performance computing in a rugged 6U chassis to meet most Mil/Aero application needs. A computing-focused platform supports up to eight Intel 5600 series server-class processors and 4TB of storage. A more balanced configuration will support two compute processing blades and two storage blades with total platform storage of 8TB. Both configurations support redundant 10G switches. Support for VMware ESXi allows the use of multiple operating systems, consolidation of applications onto a single blade and the use of vSphere 4.0 for fault-tolerant configuration.

The C2 Server also offers a 33% weight reduction over equivalent rack-mount server systems because ATCA is a bladed architecture that shares a common enclosure, fan and power supply resulting in significant weight savings. A typical system will weigh 97 lbs, light enough for two people to lift.

FEATURES

› Designed for MIL-STD-810, the C2 Server's rugged 6U chassis allows the system to meet the necessary shock and vibration standards for equipment transport.

› The RadiSys C2 server is based on an LCR ruggedized 6U, 6-slot ATCA Chassis and has been validated with:
  – RadiSys ATCA-2210 switches
  – RadiSys ATCA-4500 series of processor blades including the ATCA-4500, ATCA-4550, ATCA-4555 and ATCA-4580
  – 10G iSCSI Edge Storage blade

For more information, contact: info@radisys.com

www.mil-embedded.com/p46025

Trenton

2350 Centennial Drive • Gainesville, GA, 30504 USA
+1.770.287.3100 or 1.800.875.6031
www.trentontechnology.com/products/military-computers

Trenton MIL-STD-810 & Rugged Military Computers

Does your application require the computing platform to carry a military certification such as MIL-STD-810, or is a "designed to meet" system appropriate? The answer to this question is critical because it sets the stage for all remaining system hardware decisions. Trenton military computers are available in two different classifications: MIL-STD-810 computers and rugged "designed to meet" computing solutions. Both varieties come in a wide variety of chassis form factors to meet the diverse needs of the military. Checkout Trenton's full complement of military computers at: www.trentontechnology.com/products/military-computers, or drop us an e-mail or give us a call to discuss your military computer system requirements.

FEATURES

› Trenton MIL-STD-810 Computers – TMS2700, TMS4700, TMS4701 and TMS6700
› Trenton Rugged Military Computers – Rack-mount computers that range from 2U to 5U
› All products designed and built in the U.S.A.
› Feature long-life embedded components to support extended military project life cycles.
› Five-year factory warranties are standard with Trenton hardware
› 7-10 year product availability for key Trenton system components
› Standard and custom enclosure options
› Standard and custom board design and manufacturing services
› Designed for the realities of the military computing environment

For more information, contact: jrenehan@trentontechnology.com

www.mil-embedded.com/p47439

Suntron Corp.

Corporate Headquarters: 2401 W. Grandview Rd., Phoenix, AZ 85023 USA
1.855.SUNTRON (855.786.8766)
www.suntroncorp.com

Mission-critical, ruggedized VME/VPX/ATCA/cPCI Systems

Custom Ruggedized Displays & All-in-ones

Leading supplier of custom-driven solutions to many of the world's top companies in the Defense & Aerospace sectors.

Suntron is an electronic manufacturing service provider of engineering, manufacturing and Commercial-Off-The-Shelf (COTS) custom systems solutions that support highly configurable, highly complex products. Our team of professionals is dedicated to ensuring you receive the highest quality product on time.

Suntron's System Solutions Group allows OEMs and large end users to outsource non-core functions in order to improve overall efficiency and reduce the costs associated with the design implementation and manufacturing processes.

We understand what it takes to ensure product reliability when failure is not an option.

SERVICES

> Sustained Engineering Services
> Complete EMS Offering
> Prototype & Production
> Application Load & Test
> Custom System Design
> RMA & Spares Depot
> Direct Order Fulfillment

ITAR Registered –
ISO 9001:2008 –
AS9100

Parvus Corporation

3222 Washington Street • Salt Lake City, UT 84115 USA
800-483-3152
www.parvus.com

DuraMAR 5915 – Mobile IP Router Subsystem

The DuraMAR® 5915 is a rugged Commercial Off-the-Shelf (COTS) Cisco IOS-managed mobile router integrated with Cisco's 5915 Embedded Services Router (ESR) card in an ultra-rugged chassis optimized for harsh military and civil vehicle/aircraft installations. An ideal solution for IP networking technology refresh and situational awareness applications, including those seeking a migration path for legacy Cisco 3200 (3230/3250/3270)-based router subsystems, the DuraMAR 5915 features dual WAN uplinks and is available as either a stand-alone 5-port network router or with an integrated Gigabit Ethernet switch for a total of 19 Ethernet ports.

Optimized for Size, Weight and Power (SWaP) sensitivity as well as mechanical robustness under extreme environmental conditions per MIL-STD-810G, the DuraMAR 5915 enables prime contractors and civilian agencies to deploy Cisco Mobile Ready Net capabilities, including data, video, and voice services virtually anywhere LAN or WAN connectivity may be required, especially in mobile, airborne, ground, manned or unmanned vehicle applications.

FEATURES

> Ruggedized Cisco 5915 Router with Enterprise IOS Software for Robust Security, Management, QoS, VLAN, IPv4/IPv6 Routing, IP Mobility, Interoperability

> Modular, Open Architecture Rugged COTS PC104 Hardware Design

> Integrated Parvus Gigabit Ethernet Switch Option to Scale Port Density to 19 Ethernet Ports for Higher Port Density Application Requirements

> Stand-alone Router: 5 Ethernet Ports (3x 10/100 Switched, 2x 10/100 Routed)

> Router + Switch: 19 Ethernet Ports (15 GigE Switched, 2x 10/100 Switched, 2x 10/100 Routed)

For more information, contact: sales@parvus.com

www.mil-embedded.com/p53148

Parvus Corporation

3222 Washington Street • Salt Lake City, UT 84115 USA
800-483-3152
www.parvus.com

DuraNET 3000 – Ruggedized Cisco IE-3000 Ethernet Switch

The DuraNET® 3000 is a ruggedized version of Cisco Systems' IE-3000 industrial Ethernet switch, specifically hardened for use in demanding military/civil IP networking technology refresh applications. This fully managed network switch delivers the security, advanced Quality of Service (QoS), high availability, and manageability that customers expect from Cisco IOS-based switching technology, including optional Layer 3 IP routing services. Designed with mechanical enhancements to support deployment of data, video, and voice services in extreme temperatures, shock, vibration, humidity, as well as exposure to dust, water, and EMI/EMC environments, the unit requires no active cooling, is completely sealed, and provides interfaces over MIL-C-38999-style connectors.

Leveraging the modular nature of the Cisco IE-3000 platform, Parvus' DuraNET 3000 subsystem comes in several configurations to flexibly scale Ethernet port density from ten (10) to eighteen (18) or twenty-six (26) Ethernet ports.

FEATURES

> Ruggedized Version of Cisco IE-3000 Industrial Ethernet Switch with Best-in-Class Layer 2 LAN IOS or Layer 3 IP Services IOS Software Support

> Optional Integrated IE-3000 Switch Expansion Module(s) to Scale Port Density from 10 to 26 Ethernet Ports to Meet Application Requirements

> Designed to meet MIL-STD-810 Shock, Vibration, Thermal Conditions

> >-40 to +71C Fanless Extended Temperature Operation with No Moving Parts

> 2x Gigabit Ethernet Uplink Ports, plus 8x, 16x, or 24x Fast Ethernet Ports

For more information, contact: sales@parvus.com

www.mil-embedded.com/p53149

Sensors and RF: Digital

OceanServer Technology, Inc
151 Martine Street • Fall River, MA 02723 USA
508-678-0550
www.ocean-server.com

OS4000-T Three Axis Digital Compass

The OS4000-T is an ultra-small and highly accurate OEM compass used for embedded applications. The tilt-compensated electronic compass is packed into a 0.6" SQ. PCB weighs only 1 gm and features 3-axis magnetic sensors and 3-axis accelerometers. Offered in a through-hole package, the device is designed for mounting on a system board and talking via a TTL interface for a wide range of uses including antenna positioning, heading requirements and various other applications.

OS4000-T nano compasses are priced at $249.00 ea. qty 1, $89.74 ea. qty 500, with larger quantity discounts offered. Developer's kits are $399.00.

OceanServer
T E C H N O L O G Y , I N C

FEATURES

› Compass accuracy: 0.5 degrees RMS level heading
› Tilt-compensated (electronically gimbaled)
› Hard- and soft-iron compensation routines
› High data update rate to 40Hz
› 50 MIPS processor supporting IEEE floating point math
› 24-bit differential Analog to Digital converters from Analog Devices
› ASCII sentence output in several formats, NMEA checksum

For more information, contact: sales@ocean-server.com   **www.mil-embedded.com/p53134**

## Sensors and RF: Wireless

octoScope
225 Cedar Hill Street, Suite 200 • Marlborough, MA 01752 USA
978-222-3114
www.octoscope.com

octoBox small anechoic test chamber

octoBox enables controlled over-the-air (OTA) testing of conventional and MIMO wireless devices in a customizable refrigerator-sized anechoic (non-echoing) enclosure that replaces expensive walk-in anechoic chambers and facilitates test setup. Offering repeatable and accurate test methodology, octoBox cuts development and production costs for multi-mode wireless devices and speeds up time-to-market by several months.

As the world becomes more connected, multiple radios (2G/3G, LTE, Bluetooth, Wi-Fi, GPS) are incorporated into a variety of products. Devices such as smartphones, pads, and military radios with multiple integrated antennas require rigorous OTA testing. Existing solutions are either too small to accommodate far-field conditions that make OTA testing accurate, or too large, rendering them impractical for production and burdensome for R&D. The octoBox test platform offers OTA test stability and repeatability in an optimum form factor that fits easily in a production lab or an engineer's office in an area no larger than a typical desk.

octoScope
RF and Wireless Technology Experts
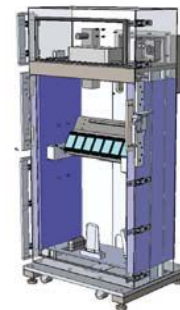
FEATURES

› octoBox is a dual-chamber RF enclosure that houses isolated and well-controlled over-the-air (OTA) test setup for production, QA and R&D test. Incorporating test instrumentation, test antennas, and an easily accessible chamber for the device under test (DUT), octoBox is a space-efficient, self-contained test station.

› The Master chamber at the top typically houses test equipment or a partner device transmitting to and receiving from the DUT. The DUT chamber houses the DUT and the test antennas.

For more information, contact: sales@octoscope.com   **www.mil-embedded.com/p53142**

Sealevel Systems, Inc.
2779 Greenville Highway, PO Box 830 • Liberty, SC 29657 USA
864-843-4343
www.sealevel.com/mes/9065

ACC-188 USB Synchronous Serial Radio Adapter

The ACC-188 USB synchronous serial radio adapter and free software from the Defense Information Systems Agency (DISA) upgrades tactical radios with the capability to send and receive IP data such as email, text messages, GPS maps, images, and coordinates.

The ACC-188 operates in conjunction with standard PDA-184 software developed by and available from DISA. The PDA-184 software provides a Graphical User Interface (GUI) that allows radio users to transmit and receive a variety of data types at much higher speeds than is possible with comparable, proprietary solutions.

A key advantage of the ACC-188 is that it enables interoperability among the various radio brands and models used by the defense community. The ACC-188 is compatible with any tactical radio that has a synchronous communication port using MIL-STD-188-184. This includes the most prevalent brands and models: Raytheon AN/PSC-5D and ARC-231, Harris AN/PRC-117F and AN/PRC-150, Thales AN/PRC-148, Rockwell Collins ARC-210, Motorola LST-5B and LST-5C.

The PDA-184 software is government-developed, government-owned by DISA, and includes these features:

• Saves thousands of dollars per radio compared to upgrading hardware and purchasing expensive proprietary software
• Interoperable (send/receive data with different radio models and brands)
• Implements MIL-STD-188-184 data waveform
• High-speed data throughput
• Easy-to-use Java-based GUI
• Runs in Microsoft Windows Vista/XP/2000

The cable and integrated assembly consists of a printed circuit board that is tested and then encapsulated using an over-mold process. This process results in a rugged, shielded, battlefield-ready product capable of withstanding harsh environments. One end of the cable includes a standard type-A USB connector, suitable for use with any USB-enabled computer. The other end of the cable includes a connector specific to a brand or model of field radio.

ARC-231
Raytheon

DB-25
Universal/DMR

PRC 117G/152
Harris

PRC-117F/150
Harris

PSC-5D
Raytheon

PRC-148
Thales

FEATURES

> Interoperable (send/receive data with different radio models and brands)

> Implements MIL-STD-188-184 data waveform

> High-speed data throughput

> Easy-to-use Java-based GUI

> Runs in Microsoft Windows Vista/XP/2000

> Overall cable length is 60"

> USB cable segment is 12", cable segment with military connector is 43.8", and plastic over-mold is 4.2" in length

> PDA-184 software is available from DISA

For more information, contact: sales@sealevel.com

Sensors and RF: Wireless

www.RFM.com (RF Monolithics, Inc.)
4441 Sigma Road • Dallas, TX 75244 USA
972-789-3824
www.RFM.com/products/wi-fi_bluetooth.php

Wi-Fi and Wi-Fi + Bluetooth Combo Modules from www.RFM.com

The WLS-series modules are ideal for OEMs wanting to quickly add Wi-Fi and/or Bluetooth standards-based connectivity within a broad range of products in audio, video, or data applications for the worldwide market.

The WLS Wi-Fi + Bluetooth combination modules are designed to fit into small spaces and are slightly smaller than a U.S. dime. They are optimized for RF performance featuring best-in-class WLAN and Bluetooth coexistence technology on a single chip from TI. The modules include high-efficiency front-end circuits plus a DC-DC converter – all delivered in a very small form factor.

Minimal external circuitry is required to complete a radio design; add an antenna, power source, processor, and associated interface hardware and the radio hardware design is complete.

The benefits to designers are low-cost, efficient design cycles and fast time-to-market.
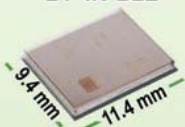
FEATURES

> Slightly smaller than a U.S. dime, taking up very little board space
> High level of integration supports efficient design cycles for faster time-to-market
> Best-in-class TI SoC – now with Bluetooth Low Energy version 4.0, 2.1 plus EDR, Power Class 1.5
> Minimal external circuitry required to complete a radio design, minimizing BOM costs
> RFM design support provided at lower volumes compared to other vendors to get you through certifications and to the market faster
> Available only from www.RFM.com, a specially designed TCXO (XTC7012) to provide years of frequency stability for high speed data and healthcare applications

For more information, contact: sales@rfm.com   **www.mil-embedded.com/p53159**

Software Enterprise: Application Life-cycle Mgt. (ALM)

MKS, a PTC Company
1815 South Meyers Road, Suite 220 • Oakbook Terrace, IL 60181 USA
800-613-7535
www.mks.com

Integrity

Integrity, a PTC product (formerly MKS Integrity), accelerates software innovation through managing the software system lifecycle management process by connecting all engineering artifacts, including requirements, models, code and test, ensuring comprehensive life-cycle traceability. Engineering teams improve productivity, which ultimately drives more innovative products into the market.

Integrity's open architecture integrates disparate tools into a stream-lined engineering process, allowing orchestration of engineering change and collaboration across the technology supply chain.

A totally unique product for companies looking to improve their competitiveness in the market, Integrity offers organizations:

• Engineering artifact orchestration
• Complete ecosystem integration
• True enterprise-wide value

FEATURES

> **Disciplines:** Integrity offers deep and rich functionality to support all of the traditional disciplines across the software system management life cycle, including: Change Control, Configuration Management, Requirements Management, Test and Quality Management, Source Code Management, Build/Release/Deployment Management, Issue/Defect Management, Process/Workflow Management, and Project/Portfolio Management.

> **Aerospace/Defense Industry:** Enables managed innovation and collaboration among large global teams (often also facing regulatory or export control scrutiny) to produce systems and products with unprecedented levels of complexity and interdependence.

For more information, contact: info@mks.com   **www.mil-embedded.com/p52830**

Annapolis Micro Systems, Inc.

190 Admiral Cochrane Drive, Suite 130 • Annapolis, MD 21401 USA
410-841-2514
www.annapmicro.com

CoreFire

Develop your application very quickly and easily with our CoreFire™ FPGA Application Builder, which transforms the FPGA development process, making it possible for theoreticians to easily and quickly build and test their algorithms on the real hardware that will be used in the field.

Use CoreFire's graphical interface to drag and drop library elements onto the design window. Modify your input and output types, numbers of bits, and other core variables by changing module parameters with pull-down menus. The modules automatically provide correct timing and clock control. Insert debug modules to report actual hardware values for hardware-in-the-loop debugging. Hit the Build button to check for errors and as-built core sizes and to build an encrypted EDIF file. Use the Xilinx ISE tool to place and route each FPGA design. Modify and use the jar file or the C program created by the CoreFire Build to load your new file into your WILDSTAR and I/O card hardware. Use the CoreFire Debugger to view and modify register and memory contents in the FPGA and to step through the dataflow of your design running in the real physical hardware.

Our extensive IP and board support libraries contain more than 1,000 proven, reusable, high-performance cores, including FIR and CIC filters, a channelizer, and the world's fastest FFT. We support conversion between data types: bit, signed and unsigned integers, single precision floating point, integer and floating point complex, and arrays. A few of the newly added array cores include array composition and decomposition; slice, parallelize, serialize, repack, split, merge, reorder, rotate, and concatenate transformations; matrix math, sliding windows, and convolutions.

The combination of our COTS hardware and CoreFire enables our customers to make massive improvements in processing speed while achieving significant savings in size, weight, power, person-hours, dollars, and calendar time to deployment.

POWER SPECTRUM COMPUTATION AND SORT FOR MAXIMUM
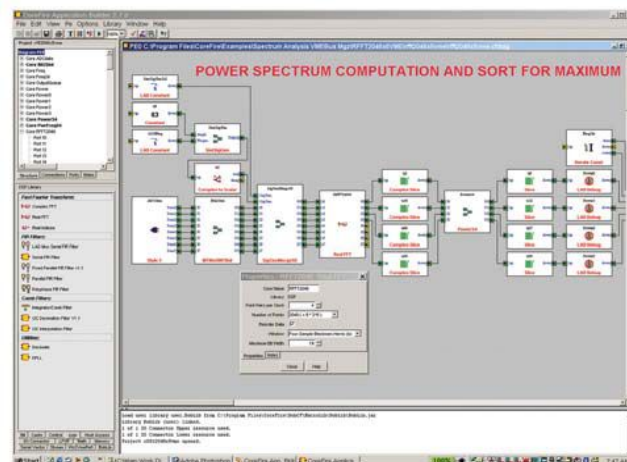
FEATURES

› Dataflow-based – automatically generates intermodule control fabric

› Drag-and-drop graphical interface

› Work at high conceptual level – concentrate on solving algorithmic problems

› Hardware-in-the-loop debugging

› More than 1,000 modules incorporate years of application experience

› Reduce risk with COTS boards and software

› Save time to market

› Save development dollars

› Easily port completed applications to new technology chips and boards

› Training and custom application development available

› Achieve world-class performance; WILD solutions outperform the competition

› Annual node locked or networked license; includes customer support and updates

---

For more information, contact: wfinfo@annapmicro.com   **www.mil-embedded.com/p33544**

## Software/Middleware: Dev. Environment/Tools

Impulse Accelerated Technologies

800 Bellevue Way NE, Suite 400 • Bellevue, WA 98004 USA
425-605-9543
www.ImpulseC.com

Impulse C-to-FPGA Optimizing Compiler

Accelerate image, signal and data processing by moving C algorithms to FPGA. Impulse C provides a fast path to FPGA hardware with high-level synthesis based on ANSI C. Impulse C enables highly iterative, software-oriented design methods and works with many FPGA-based platforms. Impulse C is the most widely used C-to-FPGA software. Join NASA, DoD, Army, Navy, Air Force, BAE, JPL, LANL, LLNL, ORNL, Northrop Grumman, Raytheon, Lockheed Martin, Harvard, MIT, Toyota, Nissan, Fujitsu and hundreds of others to design:

- **Image Processing and Object Capture** – Offload image filtering computations for real-time, streaming video processing.

- **Digital Signal Processing** – Generate heavily pipelined, hardware-accelerated DSP applications.

- **Embedded Processor Acceleration** – Speed embedded computations by 100X or more.

- **Low-latency Network Packet Sniffing** – Create server-based applications for complex data searches, filtering and simulations.
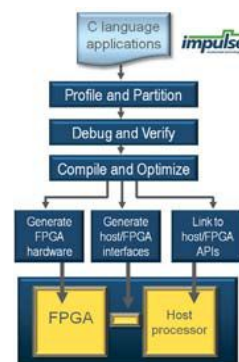
FEATURES

› Profile and partition between the processor and the FPGA accelerator. Use multiple parallel processes for increased performance.
› Verify and debug using familiar C-language development tools. Use the Impulse Application Monitor to analyze parallel dataflow.
› Compile and optimize using interactive, graphical tools. Automatically parallelize and pipeline your critical C code.
› Automatically generate FPGA hardware ready for use with your selected FPGA-based platform.
› One-key export of VHDL files to ModelSim.
› Automatically generate host FPGA, memory, I/O and other interfaces for FPGAs and FPGA-based development boards.

For more information, contact: info@ImpulseC.com   **www.mil-embedded.com/p53154**

Software/Middleware: RTOS and Tools

Esterel Technologies

100 View Street, Suite 208 • Mountain View, CA 94041 USA
Ph 1-650-641-3250 • Fax 1-650-988-8541
**www.esterel-technologies.com**

The SCADE® Product Family

Esterel Technologies' SCADE® product family provides complete solutions for developers of critical systems and software applications. All SCADE solutions easily integrate, allowing for development optimization and increased communication among team members.

SCADE Suite and Display Code Generators have been qualified/certified at the highest level of safety across six market segments by more than ten safety authorities worldwide, including DO-178B up to Level A – Aerospace and Defense Applications by FAA, EASA, Transport Canada and ANAC; IEC 61508 up to SIL3 – Transportation and Industrial Applications by TÜV SÜD; EN 50128 up to SIL3/4 – Rail Transportation Applications by TÜV SÜD, EBA and Certifer.

SCADE users report the following development and verification cost improvements:

- Aligns the design process according to safety standard objectives
- Reduces development costs by **50** percent on average
- Speeds certification so that it is achieved two times faster
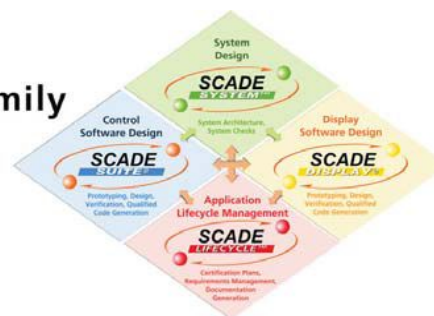
FEATURES

› SCADE Suite for control logic and algorithm development
› SCADE Display for display and HMI development
› SCADE System for system architecture design
› SCADE LifeCycle for application lifecycle management
› SCADE Solutions for ARINC 661-compliant systems for development of ARINC 661-compliant avionics displays and user applications.

For more information, contact: sales@esterel-technologies.com   **www.mil-embedded.com/p53143**

LynuxWorks, Inc.

855 Embedded Way • San Jose, CA 95138 USA
800-255-5969
www.lynuxworks.com/virtualization/hypervisor.php

LynxSecure 5.0

This new release of the award-winning LynxSecure separation kernel and hypervisor adds significant performance increases for fully virtualized guest operating systems (OSs) by utilizing new hardware technologies and offering 64-bit and SMP guest OS virtualization support.

LynxSecure 5.0 adds a secure device sharing facility for systems with limited physical devices that complements the existing direct device assignment capability. The same highly secure virtualization solution used in safety- and security-critical military, medical and avionics embedded applications can now also be used in Enterprise systems to support secure client virtualization, multi-tenancy and hardware appliances.

LynxSecure conforms to the Multiple Independent Levels of Security/Safety (MILS) architecture and has been designed to be certified/evaluated to the highest levels (Common Criteria EAL 7, TSABI, SABI PL-5). LynxSecure allows multiple "guest" operating systems to run in their own secure partitions. Guest OSs include the LynxOS family, Linux, Windows, Solaris, and more.
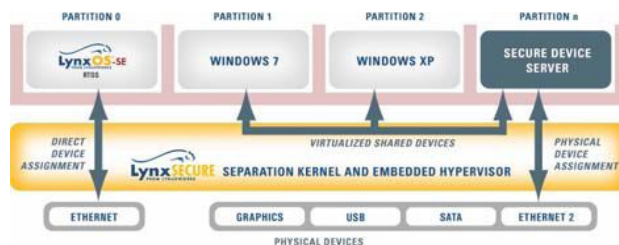
FEATURES

> Optimal security and safety – the only operating system designed to support CC EAL 7 and DO-178B Level A

> Virtualization technology - supports multiple heterogeneous operating system environments on the same physical hardware using Intel VT hardware

> Highly scalable - supports Symmetric Multiprocessing (SMP) and 64-bit addressing for high-end scalability

> Support for open standards - supports 100% binary compatibility for Linux or POSIX-based software applications to migrate to a highly robust, secure environment

> Support for latest Intel Quad Core i7 Sandy Bridge processors

For more information, contact: inside@lnxw.com   **www.mil-embedded.com/p47745**

Highland Technology, Inc.

18 Otis Street • San Francisco, CA 94103 USA
415-551-1700
www.highlandtechnology.com/DSS/V420DS.htm

V420 Resistance Simulator Module

The V420 is an eight-channel, isolated, programmable resistance simulator. It uses entirely solid-state simulation to eliminate transient errors associated with relay switching, and to provide for monotonic resistance changes. Channels are individually programmable in four resistance ranges, from 5 ohms to 65.5 Kohms. Channels can operate from microvolts up to 35 volts/50 mA. Overload protection and channel-error detection are included.

Any channel may be rerouted, under software control, to the front-panel calibration-check connector, allowing the board to be tested for calibration by a precision DVM without disturbing the connections to the equipment under test. The BIST option adds an onboard precision ohmmeter, allowing loopback checking and full automatic self-test. Optional external termination and bridge-completion boards are available.

FEATURES

> Eight-channel, isolated, wide-range, 2/4 wire resistance simulator.

> Fast, monotonic, glitch-free resistance value programming.

> Suitable for simulating resistors, platinum/copper RTDs, thermistors, strain gages, and load cells.

> Users may directly program resistances or RTD simulation (copper or platinum) temperatures.

> Four programmable resistance ranges from 5 ohms to 65.5 Kohms.

> Any channel is switchable to front-panel calibration-check connector for in-system test without disconnecting field wiring.

> Simple register-based VME interface with no handshaking required.

> Optional Built-In Self-Test (BIST) and mil/COTS versions available.

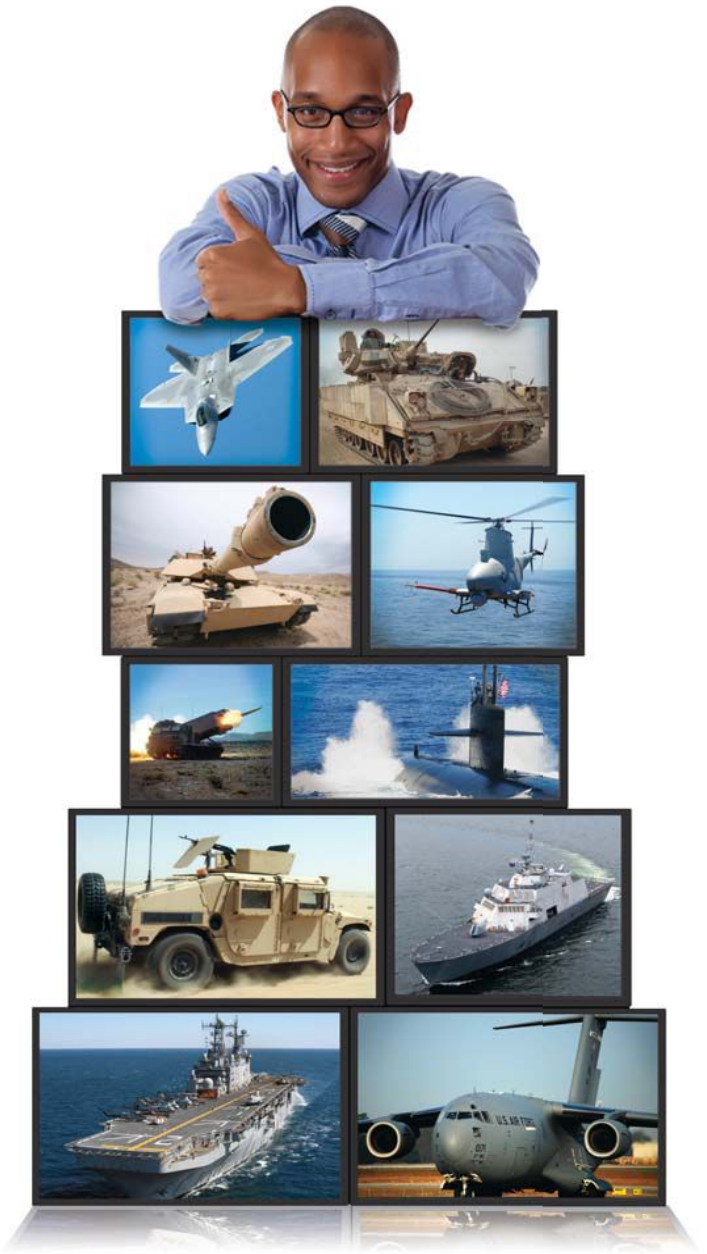> Optional termination panels and bridge completion boards available.

For more information, contact: info@highlandtechnology.com   **www.mil-embedded.com/p41872**

SPEED YOUR TIME TO MARKET

Signal & Image Processing
FPGAs and FMCs

CHAMP-FX3

FMC-518

Radar Processing
DSPs and GPGPUs

CHAMP-AV8

VPX6-490
Just Announced

Flight Control
Multi-Platform Mission Computers

MPMC-9350 & MPMC-9310 (VPX, VME or CompactPCI®)

Do you need innovative solutions or complete integration services for high-density data processing with optimized I/O? Do you face challenging size, weight and power (SWaP) constraints?

From board-level products to fully integrated systems, we provide leading-edge OpenVPX™ solutions for the most rugged operating conditions. Reduce your risk, speed your time to market and lower your development costs with our COTS products, mission computers, or modified COTS (MCOTS) engineering services.

Data Management
Rugged NAND Flash Storage
XMC/PMC-550

System Connectivity
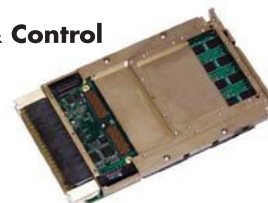Secure Ethernet Routers
VPX3-685 Fireblade

Command & Control
Single Board Computers
VPX3-1256

CURTISS WRIGHT Controls
Embedded Computing

sales@cwcembedded.com   cwcembedded.com