Extension MEDIA

Engineers' Guide to Industrial Computing

Bring Legacy Storage Interfaces into the Modern World

PC/104: What's Old Is New Again

Meet the Challenge of Securing the Smart Grid

The Industrial Computer Afterlife

www.eecatalog.com

EECatalog

Gold Sponsor

MCN

ISO 9001:2008 & AS9100C CERTIFIED

ITAR REGISTERED

VECTOR

ELECTRONICS & TECHNOLOGY, INC.
A FINE TECHNOLOGY GROUP

Since 1947

MADE IN THE USA

VME / VXS / cPCI® Chassis, Backplanes & Accessories

NEW!

Superior cooling with push-pull blowers 3U and 6U Vertical Cards

Chassis and Rack Accessories

1, 2 and 3-slot PICMG 2.11 Power Backplanes

Mil-1-46058-C Conformal Coating Available for all VECTOR backplanes

Hi-speed VITA ANSI/VITA 1.1-1997 monolithic backplanes (Hi Current VITA 1.7 compliant) with Electronic Bus-Grant (EBG), Surface mount devices, fully tested and certified. MADE in USA, ships in 2-3 days

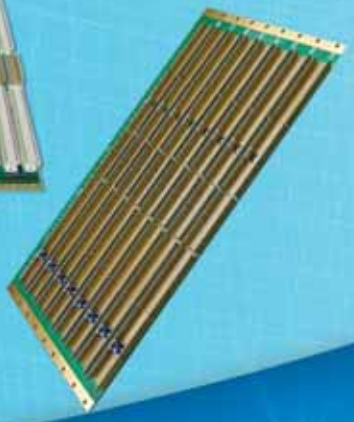

(800)423-5659

WWW.VECTORELECT.COM

MADE IN U.S.A.

Welcome to the 2014 Engineers' Guide to Industrial Computing

To subscribe to our series of Engineers' Guides for embedded developers and engineers, visit:

www.eecatalog.com/subscribe

EECatalog

Engineers' Guide to Industrial Computing 2014

www.eecatalog.com/industrial-computing

Vice President and Publisher

Clair Bright
cbright@extensionmedia.com
(415) 255-0390 ext. 15

Editorial

Editor-in-Chief
Chris A. Ciufo
cciufo@extensionmedia.com

Managing Editor
Cheryl Berglund Coupé
ccoupe@extensionmedia.com

Editorial Director
John Blyler
jblyler@extensionmedia.com
(503) 614-1082

Creative/Production

Production Manager
Spryte Heithecker

Media Coordinator
Yishian Yao

Graphic Designers
Nicky Jacobson
Caldin Seides

Senior Web Developer
Mariam Moattari

Advertising/Reprint Sales

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

Sales Manager
Michael Cloward
mcloward@extensionmedia.com
(415) 255-0390 ext. 17

Marketing/Circulation

Jenna Johnson

To Subscribe

www.eecatalog.com/subscribe

Extension MEDIA

Extension Media, LLC
Corporate Office

President and Publisher
Vince Ridley
vridley@extensionmedia.com

Vice President, Sales
Embedded Electronics Media Group
Clair Bright
cbright@extensionmedia.com

Vice President, Business Development
Melissa Sterling
msterling@extensionmedia.com
(415) 970-1910

Special Thanks to Our Sponsor

MEN

Contents

Meet the Challenge of Securing the Smart Grid

The Internet of Things (IoT) promises to bring greater efficiency and functionality to smart grid and industrial applications, but as this infrastructure is increasingly exposed over the Internet, new security challenges arise as well.

Tony Magallanez, McAfee, part of Intel Security

The Internet of Things (IoT) promises to bring greater efficiency and functionality to smart grid and other energy applications by enabling operators to more effectively monitor and manage both capacity and demand. However, as this infrastructure is increasingly exposed over the Internet, new security challenges arise as well. These challenges apply not just to smart energy and utility systems but to other "Big Data" applications such as industrial and manufacturing that also have extensive infrastructures.

Data plays an important role in the day-to-day business decisions of managing a smart grid deployment. Specifically, the movement of energy is based on current capacity, availability and demand information that need to be collected, analyzed and acted upon in real-time. Devices across the entire energy grid are continuously in communication with each other, analyzing data and making decisions based on this data. The greater the ability to correlate this data across large numbers of devices, the higher the efficiency of the overall operation.

The aspect that makes the smart grid so useful—it's interconnectedness—is what also makes it most vulnerable. Every connected device is another entry point into the overall network and affords an opportunity for tampering. Many of these devices are also connected via the Internet, providing direct access to anyone with the device's IP address. Furthermore, a breach into one part of the network can expose every other part of the network.

One of the primary security risks is that the smart grid is connected to physical assets like generators whose failure can cause considerable damage and injury. Consider Stuxnet, the worm that attacked several energy facilities around the world, including Iran's nuclear enrichment infrastructure. A study conducted by McAfee and the Center for Strategic and International Studies (CSIS) at the time (http://www.mcafee.com/us/resources/reports/rp-critical-infrastructure-protection.pdf) discovered that nearly half of the respondents in the energy sector reported that they had found Stuxnet on their systems. The cost of downtime as the result of cyberattacks on critical infrastructure averaged US $6 million per day. The study also found that one in four power companies globally had been the target of extortion related to cyber threats, resulting in losses estimated as high as hundreds of millions of dollars.

Not all threats are on the large scale of a Stuxnet or Duqu. For example, localized attacks on individual energy meters pose a serious risk to utility companies. Imagine malware posted on the Internet that allowed consumers to hack their meters and reduce their recorded usage in a way that was difficult to detect. Direct revenues losses across the grid could be tremendous.

Note that the objective of a cyberattack may not be to cause damage. The objective of the "Night Dragon" worm, for example, appears to be the collection of sensitive information and system weaknesses that could be exploited at a later time.

Smart Grid Vulnerabilities

To understand the scope of the security issues energy and utility companies face, consider that an estimated 70 percent of the existing energy grid is more than 30 years old. Given that security within the energy infrastructure has only become a concern in recent years, much of today's energy infrastructure has been designed with availability as the prime consideration. Further compounding the challenge is the fact that energy and utility companies do not have the internal expertise they need to secure their deployments. Thus, this infrastructure is especially vulnerable.

The magnitude of any Big Data application increases the necessity for comprehensive security. To be effective, security needs to include not just protecting data but ensuring that systems execute authorized applications only. In addition, security needs to extend out to the edge, including both remote facilities and endpoint devices like utility meters. Another important issue is that as smart grid devices are increasingly built using off-the-shelf components, they become increasing generic as well. Thus, if a vulnerability in a device is found, the vulnerability will be system-wide.

Because the smart grid is so dependent upon intelligent endpoint devices, these embedded systems are a primary cyber target. It is essential, therefore, to integrate security

solutions natively in these devices. Security also needs to be addressed as a cohesive solution across each endpoint and the network, protecting not just data security but code integrity as well.

Numerous security technologies are available to protect endpoint devices. Because of the threat to PCs, many people are already aware of the value of advanced encryption, firewalls and antivirus and anti-malware protection. The complexity, exposure and vulnerability of the energy infrastructure, however, requires a much more robust security strategy.

Implementing Effective Security

One method used to attack smart grid devices is malware. Malware is malicious software that a cyberattacker attempts to load onto a device to subvert its operation. Malware can reside quietly on a system, giving no indication of its presence until an attack is launched.

To protect against malware, systems can employ application blacklisting. Blacklisting works by identifying the signature of malicious code within a system and blocking any attempt to upload or execute code with such a signature. Advanced blacklisting techniques also make use of real-time analytics to extend protection to include the behavior of an application or file in determining whether software is potentially malicious. By itself, however, blacklisting is not enough to thwart malware attacks as it can take hours or even days to detect new malware and update systems with the appropriate signature.

Whitelisting increases malware security by ensuring that an embedded device only accepts commands or updates from a trusted application. Whereas blacklisting identifies software that cannot be executed, whitelisting limits execution only to explicitly enabled applications. The result is that if malware does succeed in infiltrating a device through its system interfaces, any attempts to take over the system or broadcast sensitive information will be ignored. In addition, the intrusion will be reported upon detection and exposed.

To simplify creating the approved application list, active whitelisting is available which employs a "trust but verify" model that validates and approves updates before installing them. For even greater protection, the approved configuration checklist can be duplicated in another location, enabling a system to verify the integrity of its checklist before taking action.

Among the most advanced technologies available for use today is Security Information Event Management (SIEM). SIEM technology provides analysis of security event data across the network so operators can understand trends and identify anomalies. The ability to assess behavior at the network level enables real-time visibility into operations that has never been possible before. For example, trends can be identified as they develop, enabling operators to evaluate them and respond immediately.

To experience the full benefits of SIEM, however, the underlying architecture has to be able to support the massive data volume associated with monitoring and analyzing data from devices throughout the smart grid. The McAfee Enterprise Security Manager (ESM) appliance, for example, is rated to handle 50,000 security events per second. This enables rapid analysis of security event data, reducing the processing time from days and hours to minutes. This not only increases the efficiency with which a network can be managed, it enables enterprise-class security that can quickly assess, identify and resolve threats before they can spread and lead to system failure. It also enables compliance reports to be compiled just as quickly. Futhermore, when SIEM is integrated with other security technologies, the ability of the system to assess vulnerabilities is further enhanced by mapping against factors such as confidentiality, integrity and availability.

Many security features are implemented in software and operate above the operating system. As a result, systems are still vulnerable to ever more sophisticated rootkits designed to hide the presence of malware and enable deep penetration into a targeted system. To provide protection that works below the operating system, hardware-assisted technology is required. An example of such advanced security technology is McAfee DeepSAFE, which was developed jointly by Intel and McAfee. The hardware component of DeepSAFE operates independently of the operating system and as a result cannot be directly corrupted by malware. This serves to strengthen overall security within the device by hardening it deep in the computing stack or below the rootkits.

Continuous Compliance

To verify the security of a system, audits can be conducted regularly to confirm compliance of the system to established security specifications. However, compliance does not ensure security over time because compliance testing typically focuses on the audit, a single moment in time. Cyber terrorists continuously leverage technological advances to discover new ways of infiltrating networks. Compliance alone, then, cannot guarantee that security measures will never be breached.

Security can be likened to a system's immunity system. As viruses and malware evolve over time, so the immunity system needs to be able to build up resistance to address these new threats. If devices deployed today are fixed in their security implementation, how will they be able to provide sufficient security in even just a few years' time?

A robust security strategy, especially one for smart grids applications where equipment may be deployed for decades,

must be able to dynamically address changing threats. To do so, it must:

- Minimize security risks by controlling code that executes on embedded devices
- Protect the integrity of all data on devices, both from malicious modification and unauthorized access
- Manage software change control so that only authorized software can be installed on deployed devices
- Assure compliance of devices and that systems are always audit-ready
- Provide effective change policy enforcement

An effective security strategy will provide these requirements, as well as offer:

- An efficient implementation with a small footprint and low overhead
- Transparent and application independent deployment so it can be "deployed and forgotten"
- Real-time visibility into system behavior
- Closed-loop reconciliation
- Comprehensive auditing capabilities
- Intelligent threat awareness and analysis
- Streamlined device management
- loss prevention
- Support for fixed-function and legacy systems

Every approach to security has its advantages and disadvantages. By integrating these different technologies with each other, the most effective security can be achieved. For example, blacklisting and whitelisting provide protection from different sides, so to speak, by trusting "good" software and blocking "bad" software. Where one technology is vulnerable, another is strong. In addition, the effectiveness of these technologies working together is greater than the sum of their parts, such as how SIEM technology can provide more comprehensive visibility when it has access to other security components in the system.

Good security is about more than just the bottom line. Certainly, the cost of replacing damaged generators far exceeds the investment required to protect them. However, a much greater concern is our growing dependence upon the smart energy grid. By providing effective protection throughout the energy infrastructure, we can mitigate the risks that arise when devices are connected and instead enjoy the benefits and flexibility we gain with peace of mind. For more information on protecting critical infrastructure, please visit http://www.mcafee.com/us/solutions/critical-infrastructure/critical-infrastructure.aspx or http://www.mcafee.com/embedded

With more than 20 years of IT and network security experience, Tony Magallanez has worked with emerging technologies such as data and network encryption and wireless security. He has followed the evolution of computer security technology from simple file viruses to today's fast moving hybrid viruses, rootkits and advances persistent threats. A frequently quoted source among IT publications, Mr. Magallanez now assists global companies with critical infrastructure security projects. Mr. Magallanez holds a BS in electrical engineering from New Mexico State University.

Bring Legacy Storage Interfaces into the Modern World

Storage bridging and new drive technologies can extend the lifetime of storage system interfaces, increase reliability and add security features while reducing system size, weight and power.

By Steve Gudknecht, Elma Electronic

Just gimme that old-time SCSI interface. Such is the call from many folks in the embedded industry who have very reliable—and heavily relied upon—systems that still use a SCSI interface and drive technology. Ranging in applications from ship-board weapons control to semiconductor manufacturing, thousands of systems designed years ago, and in operation today, still rely on SCSI hard drives for their storage subsystems and will continue to do so for the foreseeable future.

The problem is that SCSI drive supplies are drying up. More and more hard drive manufacturers are dropping out of the business so the availability horizon is uncertain. Making investments to improve drive reliability and performance—improvements that would be beneficial even in older systems—are not always applied to a technology that is widely viewed as past its prime.

SCSI, small computer system interface, is a standard describing the physical and electrical characteristics for a parallel-attached, multi-drop, computer-to-computer or computer-to-peripheral interface. Though originally intended for connecting many types of peripheral devices, including printers and scanners, the interface has evolved over time to be almost exclusively applied to storage devices. With the architectural shift from parallel to serial I/O in embedded systems, any device that is natively connected via parallel interface is on the down side of the availability slope.

Coping with Change

In the age of shrinking budgets, greater system performance expectations and reduced downtime requirements, cost-efficient solutions are in high demand. Making the changeover from SCSI to SATA, or even PATA, drives entails hardware redesign to accommodate the new interface and its accompanying connector and form factor types.

SCSI connector pin counts range from 50 to 80, while SATA connectors contain just 7 pins, forcing board design changes to accommodate the new connector. Many currently deployed SCSI drives are 3.5 inch while SATA drives with few exceptions are 2.5 inch or smaller.

When every dollar is measured no matter the application, buying a boat load of drives for stock pile as a hedge against impending scarcity is often not an option. The real

Evolving Storage Technology - The Challenges:

- Identify suitable SATA drives to replace the old PATA or SCSI drives
- Maintain the legacy interface to the host system
- Maintain the legacy board form factor – VME, cPCI etc. – on which the drive(s) are mounted
- Maintain the legacy operating system
- Support removability where necessary
- Insert new capabilities

Table 1

heartache occurs, however, when changes to the operating system and device drivers necessary to support the new host interface are considered because new hardware often means new software.

Software changes carry system-wide implications that impact other processes and communication subsystems that end users simply don't want to touch. As a result there is an emerging need for seamless methods that bridge legacy systems to new storage technologies without fixing what isn't broke and with only the slightest impact to the budget.

Conquering the Great Divide

Bridging solutions allow the use of modern storage drives to be connected via the legacy host storage interface using form, fit and functional board-level replacement products designed to insert into specific applications. SCSI-to-SATA bridges form elegant solutions that present the legacy interface to the host, while masking any indication of the underlying SATA technology and thereby eliminating the need for software changes.

Figure 1. Dual SCSI interfaces support one wide SCSI-to-SATA and one Ultra320 SCSI-to-SATA II connection on a VME board where the original design consists of dual straight SCSI connectivity.

As you may guess, many legacy systems are Eurocard-based and for the most part specifically VME or its derivatives (with the exception of VPX) along with some cPCI. Bridging solutions can be applied to any board size, however, and modular solutions make it easy to adapt bridges mechanically and electrically for fast development cycles.

Solutions have been developed that use current devices such as SSDs, compact flash and CFAST media to replace large rotating hard drives, and tape drives that support removable cartridges. Board-level storage product suppliers must cover a short—but very important¬—list of considerations when solving a legacy storage bridging problem. (See Table 1.)

Creative Designs Ease the Transition

Embedded board designs aimed at solving this challenge use advanced bridging devices to create custom solutions that address it in several ways. Relatively simple designs consist of dual SCSI interfaces with non-removable drives. An example is a board that supports one wide SCSI-to-SATA and one Ultra320 SCSI-to-SATA II connection captured on a VME board where the original design consists of dual straight SCSI connectivity. (See Figure 1.)

More advanced conversion challenges include interfaces where the legacy hardware includes not only SCSI hard drives but also removable media like DAT drives attached via SCSI. Where removability is a carry-forward requirement, a SCSI-to-ATAPI bridging solution is necessary. ATA Packet Interface is a protocol that transports packetized SCSI commands and also carries additional capabilities including a media eject command.

DAT tape drives, much like rotating hard drives, are inherently failure-prone in relation to rugged and lightweight solid state alternatives, so given the chance, users should opt for making the solid state transition. Considerations in this case must be given to the data offload method, since removable drives must be compatible with hardware at the data offload point. (See Figure 2.)

Many older deployed systems will remain unchanged throughout their usable life and so a fixed bridging solution—such as those described above—will do the trick. But in cases where a transition to a new host system is planned, it may be beneficial and cost-effective to address the needs of both the old and the new systems in a single configurable manner.

Moving to a new system architecture is a staged and carefully managed transition that may require a significant period of overlap where both system types are deployed and supported for years. To cover this transition period, creative board designers make use of simple switching devices embedded on the doorstep of the drive that enable either a pass through of the SCSI command set or the SATA command set.

Front or rear panels would include I/O connectors for both the legacy interface and for the new SATA interface, allowing the single design to be applicable to both systems. In each case described above, custom transition modules can be designed to support rear I/O where the need arises. The dual-purpose board may be deployed as part of either the legacy system or the new system, reducing spares requirements with support for both interface options.

Interface conversion applies not only to cases involving SCSI-to-SATA but also SCSI-to-PATA and finally PATA-to-SATA and the basic framework of the solutions remains similar with the exception of the specific bridge device used. SCSI-to-PATA bridging solutions although easy to implement, are not recommended since PATA drives, like SCSI and other parallel I/O devices, are also on the down-side of the usage slope with fewer and fewer suppliers in the business.

Figure 2. Removable drives must be compatible with hardware at the data offload point.

Interface speed mismatches are inevitable, with the overriding requirement being that the bridged solution does not carry with it a reduction in data transfer rates. Many older storage subsystems support data rates lower than those attainable using today's interfaces and drives, so data-flow constraints are determined by the legacy interface. Currently embedded storage products make use of Ultra 320 SCSI-to-SATA II bridge adapter solutions, which match up theoretical bandwidths at around 3 Gb/sec...enough bandwidth to support most entrenched systems.

A Chance for Improved Performance

Beyond resolving the immediate need to keep your system obsolescence-proof into the future, there are benefits that you may take advantage of when upgrading to the latest drive technology.

Rotating hard drives make up the vast majority of all drives found in older systems. They also have the dubious distinction of being the only wear item, the least reliable and the most delicate component in the system, accounting for most of the failures and repair costs in deployed systems. With that in mind, any transition in drive technology is also the chance to move to more rugged solid state drives, thereby increasing overall system mean time between failure (MTBF) and improving long-term repair costs and environmental toughness.

Though more expensive—low-cost MLC solid state drives cost around eight times more than rotating hard drives on a per gigabyte basis—total cost of ownership approaches parity when considering improvements in equipment up time and reduced replacement costs.

Data security is more important today as compared to when most legacy equipment went into service. In addressing that new reality, solid state drives now offer multiple levels of secure erasure plus write protection and data encryption and these options can easily be accommodated when designing replacement boards. Power-hungry rotating hard drives dissipate roughly three times the wattage as SSDs and weigh up to eight times more.

Since many legacy SCSI systems use 3.5 inch drives, another way to reduce weight results from transitioning to 2.5 inch drives—either rotating or solid state. In this case a complete solution would consider any mechanical mounting changes necessary to support the smaller drive form factor as it is nested on the board.

Conclusion

Storage solutions can be replaced transparently, while maintaining backwards compatibility with operating system software and associated drivers. Implementing storage bridging technology and replacing SCSI and PATA drives with cutting-edge SATA drives can extend the lifetime of the storage system interface by five or more years.

End users will find additional immediate benefits in making the jump to SATA drives by taking advantage of the opportunity to move to more reliable solid state versions and adding security features while possibly reducing system size, weight and power. Embedded storage manufacturers, such as Elma, with a long history of innovative designs spanning multiple interface variants are in the best position to address legacy storage configuration requirements.

Steve Gudknecht is product marketing manager at Elma Electronic. He has held positions in field applications and marketing in high technology industries for nearly three decades. Steve's responsibilities include product development, product marketing, training and sales support.

Multicore Comes of Age

The move to multicore is now well on its way, in applications from smartphones to networking equipment, and the door is even cracking open for safety-critical applications.

By Cheryl Coupé, Managing Editor

Multicore is becoming an expectation in nearly every type of embedded system. While hurdles still exist in the most stringent safety-critical arenas, the technology is beginning to make inroads even there. Our roundtable panel addresses the ongoing challenges for developers and new specifications and tools to address them, as well as exciting new developments in virtualization and high-speed interconnects Participants in this roundtable are Bill Graham, director of product marketing for Wind River; Pekka Varis, CTO for ARM and DSP Processors at Texas Instruments; and Mark Throndson, director of business development at Imagination.

EECatalog: As recently as 4 years ago we had dual cores and programmers were still confused as to how to program them. Fast-forward to today's 8-core processors—how are developers coping?

Bill Graham, Wind River: Although challenges still exist with programming for the true parallelism that multicore processors bring, embedded developers are realizing that multicore processors are offering different architecture options that weren't available before. Rather than trying to reprogram their application to make use of 4, 8, 16 or more cores, they are porting their applications as-is to a single core and leveraging virtualization and the improved processing to power ratio to consolidate multiple systems. Dealing with multicore at the higher architecture level is paying off in significant ways despite the new programming challenges.

Pekka Varis, Texas Instruments: Developers for certain applications are coping just fine, but there is not one single silver bullet. Take networking for example: there are several implementations that achieve linear or nearly linear scaling with the numbers of cores. One of the more commonly known ones is 6WindGate networking stack, but similar approaches have been used before and are used in parallel. In the ARM® world, Linaro™ is standardizing this under the name OpenDataPlane (ODP). In the embedded world with a more computational focus, (for example scaling Fast Fourier Transforms (FFT)) standardization is with OpenMP on TI KeyStone multicore DSPs. However, many applications relying on heavy reuse of software programmed years ago remains a problem for developers.

Mark Throndson, Imagination: The long-standing methods of achieving higher performance through increasing clock frequency, process migration and maximizing single-thread CPU performance have yielded diminishing returns for some time. The move to multicore was a natural progression to enable performance scaling. Today, there is enough parallelism in many software workloads to readily make use of multiple cores or threads in a CPU. And given that this is a key path for increasing CPU performance, there will be a continued focus on increasing the parallelism in the software. However, thanks to the advancements in heterogeneous compute, developers today can also use the GPU for the heavy lifting part of an algorithm. Because the GPU is inherently a massively multi-threaded machine, it can handle these types of tasks much more efficiently.

EECatalog: What is the impact of the Multicore Association's new SHIM spec? What else is in the works that will improve the efficiency of multicore/many-core programming?

Graham, Wind River: Although it's an emerging specification, SHIM does promise an open standard for defining multicore architecture descriptions for software and tools. As we enter the many-core era, programming will require tool support in order to succeed; SHIM is working to make this possible. The impact of SHIM and other standards the Multicore Association is working on will create an open standards-based environment for programmers, tools developers and vendors. The Multicore Association's other working groups such as the MCAPI, are furthering this standardization effort. In a similar vein, Intel has been promoting various programming options for multicore and many-core, such as OpenMP, Intel Thread Building Blocks and most recently Cilk Plus.

Varis, Texas Instruments: Multicore Association SHIM is in an early stage. Unifying the low-level interface to allow tools and compilers to leverage what infrastructure is on a given multicore device should enable tools to focus on the bigger issues rather than SoC-specific nuances. However, for SHIM to have a significant impact it will be important that it is part of a widely adopted standard.

EECatalog: What is changing in the use of multicore processors in safety-critical applications? How are new RTOSes with separate kernel and hypervisor impacting that?

Graham, Wind River: The significant change is the acceptance of multicore-based systems by the safety certification agencies around the world. A significant hurdle is designing and developing multicore systems that meet the strict safety standards such IEC 61508 and DO-178B/C. Solutions that are offering already-proven separation technology such as ARINC-653 on multicore platforms are the most likely to succeed. Similarly, bare-metal hypervisors are also being accepted in safety-critical design. In many ways, these virtualization and partitioning technologies, already proven on single-core systems, are the key to success in the multicore transition for safety-critical systems.

Throndson, Imagination: Traditionally, ensuring reliability and security in many safety-critical applications was done by separating tasks onto multiple independent CPUs. By running an RTOS with separate kernel and hypervisor on a CPU platform, the separation and prioritization of tasks can potentially be done securely and reliably on one CPU. This can be done as a software only, or para-virtualized implementation; however using a CPU family with hardware virtualization can minimize the overhead as well as leverage the use of existing software. This doesn't mean that multicore isn't necessary; it simply means that performance increases in importance as a motivation for a multicore implementation.

EECatalog: What interesting virtualization capabilities relying on multicore are taking place?

Graham, Wind River: Multicore processors are the enabling technology for virtualization. Although completely functional on single-core systems, the improved processing to power (and size and weight) ratio means that use cases for virtualization are more compelling today. The addition of hardware support for virtualization is removing most of the overhead and real-time responsiveness and latency is now possible with embedded virtualization. Probably the most exciting thing that virtualization brings is the new architecture options embedded developers enjoy. Consolidation of multiple systems into one: for example, industrial control systems with hundreds of programmable logic controllers, user interface, data acquisition and network gateway can be integrated on a single platform. In networking infrastructure, network functions virtualization (NFV) has taken off as a way to consolidate multiple, complex and expensive custom hardware pieces into the equivalent software services running on IT, server-grade hardware. Multicore processing is the key enabler that has made virtualization a reality in these new applications.

Varis, Texas Instruments: In the embedded space, multicore devices have always needed an element of virtualization to allow sharing of peripherals and accelerators. Initial approaches relied on multiple sets of registers, one per core, or later on, one per virtual machine. But for the higher end cores such as ARM® Cortex®-A series with deep out-of-order pipelines, any accesses outside cached memory carries a performance cost.

Ring-type structures in memory coherence for the high-end cores and the I/O seem to be promising, but for simpler cores and deterministic applications, hardware queues make sense. Regardless, the hardware and associated direct memory access must be able to parse and understand the same structures as software.

Throndson, Imagination: Virtualization has traditionally been used in servers needing many nodes operating separately in parallel, but the use of this technology is expanding into an increasing variety of applications. It can also be used for mixed-mode Linux and real-time environments in separate domains, with open source application processing in one, with real-time, latency-sensitive tasks running in another, with QoS/priority.

There are many high-volume, consumer-oriented devices depending on multicore levels of performance today including smartphones, tablets, high end TVs and the list goes on. Virtualization becomes compelling as a solution in these applications to address the growing requirement for a scalable security implementation. It provides a scheme for independent, secure domains for DRM, content protection, transactions, identity protection and so on. We've recognized the importance of hardware virtualization, and that's why it's a foundational technology across the entire range of our MIPS Series5 Warrior CPUs.

EECatalog: What is the impact of inter-processor communication (IPC) busses like Intel's Quickpath or Xilinx's RocketIO or even Serial RapidIO with its RDMA? Do they extend the concept of closely-coupled multicore to non-contiguous/non-homogeneous "multiple cores"?

Graham, Wind River: These interconnected technologies promise to provide the same or better data rates that are achieved by processor busses including PCIe and fully buffered memory. This means that multicore can extend beyond a single physical processor to many processors without significant communication overhead. So yes, these extremely fast

busses will enable multicore to many of the same or different (heterogeneous multicore) types of CPUs and I/O devices. As the consolidation use case becomes more and more prevalent, combining multiple systems with multiple CPU architectures on the same circuit boards is desired. Quickpath, RocketIO and SRIO make this a reality.

Varis, Texas Instruments: Interchip busses have been used in embedded systems for a while. There are systems from dozens of processors connected with SRIO to build a base station, to supercomputers with thousands of processors. Typically the proprietary interconnect such as the two above or TI's Hyperlink, leverage high-speed SERDES and impose some restrictions to achieve more bandwidth or lower power per bandwidth than a standard interconnect. However, this approach creates fragmentation and interoperability barriers.

Serial RapidIO has a lot of good attributes, and it is a standard that has been proven to interwork, although some of the players in the market might prefer the fragmentation and success of proprietary technologies.

Cheryl Berglund Coupé is managing editor of EECatalog.com. Her articles have appeared in EE Times, Electronic Business, Microsoft Embedded Review and Windows Developer's Journal and she has developed presentations for the Embedded Systems Conference and ICSPAT. She has held a variety of production, technical marketing and writing positions within technology companies and agencies in the Northwest.

PC/104: What's Old Is New Again

Despite all the breathless excitement of the "new" Internet of Things, the PC/104 Consortium has spent its 22 years of existence helping to get the industry to this point: bringing the computer off the desktop and into the field where data is gathered and real-time decisions are made.

By Cheryl Coupé, Managing Editor

PC/104 continues to be a major player in the embedded market. We sat down with Dr. Paul Haris, president of the PC/104 Consortium, for a conversation on how PC/104 fits in with the latest industry trends. In addition to his role at the Consortium, where Paul has held positions of chairman and president for three years, as a board member for six years, and as chair of the Technical Committee for four years, he is also president and CEO of RTD Embedded Technologies, Inc. RTD has been a part of the leadership structure of the PC/104 Consortium since it help found the Consortium in the early '90s.

EECatalog: Start out by telling me what's new with the PC/104 Consortium. What's the next major challenge you'll face, and what type of specification is likely to result?

Haris: The computer market as a whole has been evolving and expanding rapidly with the advancements of technology. Since its beginning 22 years ago, the PC/104 Consortium pioneered bringing the computer to where it was needed most: out of the building and into the field where data is gathered and real-time decisions need to be made. Today's PC/104 specifications support main bus architectures and I/O capabilities. It provides expandability, upgradability and maintainability all in an inherently rugged form factor. Last year, the Consortium expanded the capabilities of the PCIe/104 and PCI/104-Express specifications to include Gen 2 and Gen 3 PCI Express speeds. The Consortium continues to closely monitor potential evolutionary paths of the embedded market and is situated to meet its challenges. As new bus architectures evolve and become industry standards, the PC/104 Consortium is ready to incorporate them in a logical and meaningful way to ensure their continued long term use and supportability.

So where will you find PC/104 in the Internet of Things? Everywhere.

EECatalog: The overall embedded market is abuzz with Intel's latest processors and ARM's latest SoCs. How do these trends affect the PC/104 set of specifications?

Haris: The stackable PC/104 architecture has been supporting both high-speed processors as well as small, embedded ones for the last 22 years. It has never been limited to any particular processor type since the stacking connector incorporates mainstream bus signaling. This is why you have seen them—whether x86, ARM or PowerPC—on PC/104 modules for many years. As processors become smaller and smaller, with ever-increasing computational power, efficiency and full-feature I/O sets, the PC/104 architecture shines by being able to create fully self-contained single board computers with the added benefit of stackable expandability for additional or specialized processor and I/O functionality. All this while keeping the overall system size to a minimum.

EECatalog: There's big growth occurring in the extended temp and harsh environment markets of automotive, transportation, mining, railway, power plants and more. How is the PC/104 Consortium addressing these rugged applications?

Haris: The PC/104 architecture has always been known as an inherently rugged architecture with its stackable bus. You will find systems in the harshest conditions of land, sea, air and space. The flexibility of having the stackable backplane on each module allows OEMs and end users to create standard and custom systems that address the thermal requirements for their particular applications. For over 20 years the transportation, mining, railway, military and other demanding industries have relied on PC/104 to meet their challenges.

EECatalog: The PC market is dying, according to myriad analyst forecasts and reported numbers by PC vendors. Yet many small form factors directly map to the PC market versus the mobile market, for example. What does this trend mean to PC/104?

Haris: It is important not to confuse the type of processor architecture with the form factor and bus architecture of many standards. What we are seeing is not a dying trend but a separation of trends. As the embedded marketplace explodes,

there is an ever-increasing need for point location of computer capabilities and the mobility of tasks and information gathering. To meet these needs, the x86 market has traditionally taken a top-down approach. While this was happening, a bottom-up emergence was occurring. Phones were getting smarter, more powerful and more efficient. But they were also based on a specialized, non-mainstream technology. As the demand for capability increased in the mobile market, so did the need for additional features and expansion capabilities such as PCI Express. In the end, these two competing processor architectures each has its unique markets, but there is much overlap in the middle. With its stackable backplane based on mainstream bus signaling, the PC/104 architecture has the capability to serve all of these architectures and markets giving maximum flexibility to the end user.

EECatalog: What are the most recent obsolescence issues faced (or being faced) by PC/104 vendors?

Haris: The PC/104 Consortium has always looked to the trends of the embedded market and it has guided its specifications accordingly. To ensure longevity, migratory paths to advanced technologies while maintaining mechanical backward compatibility have always been a top priority. This can be seen in the evolutionary progression of its stackable buses and their placements on the Consortium's 104 form factor. Manufactures and end user are thus given a timely evolution path for their past, present and future designs. This is what has given the PC/104 architecture such a prominent place as one of the longest and versatile industrial standards in the world.

EECatalog: What "play" will small form factors like PC/104 have in the movement known as the Internet of Things?

Haris: The "Internet of Things" terminology has often been thrown around very loosely. Simply, it is the world where devices are hooked through an Internet-like structure. It can span anywhere from the very small, such as discrete sensors and appliances, to the very large, like bulldozers and factory floors. It can include information gathering as well as control. It can allow instantaneous connectivity to your infrastructure wherever you are. But it is also fraught with a myriad of security risks. The PC/104 architecture has always been a part of this movement, before it was even considered a movement. One of the main points of PC/104 is the ability to distribute information gathering and computational capabilities out of the building to where it is needed most: the device. But unless it is to operate autonomously without monitoring, connectivity has been required, often through the Internet. In addition, the versatility of the PC/104 architecture has led to the creation of infrastructure devices such as firewalls, routers and switches. So where will you find PC/104 in the Internet of Things? Everywhere.

Cheryl Berglund Coupé is managing editor of EECatalog.com. Her articles have appeared in EE Times, Electronic Business, Microsoft Embedded Review and Windows Developer's Journal and she has developed presentations for the Embedded Systems Conference and ICSPAT. She has held a variety of production, technical marketing and writing positions within technology companies and agencies in the Northwest.

Convergence and Security will Drive Internet-of-Things Proliferation

By Jonah McLeod, Corporate Marketing Communications Director, Kilopass Technology Inc.

Tony Massimini has been digging into the Internet of Things (IoT) and has come up with some interesting findings. Semico Research released two reports in January this year "What Does the Internet of Things Need to Grow?" and "The Internet of Things, Augmented Reality, and Sensor Fusion," detailing what he has learned. You can also get the latest at the SemiCo IMPACT Event on April 23rd at the Biltmore Hotel in Santa Clara, California. In describing the problems confronting this potentially huge market opportunity everyone keeps referring to as the IoT, he cited a lack of unifying platform to bring a number of divergent solutions together and security as the two obstacles that need to be hurdled.

The IoT is actually a collection of siloed solutions: industrial control, personal electronics, home automation, etc., he noted. For example, industrial control, which began as a wired solution to link equipment for food, plastic, or metal casting processing and production line conveyors, machine doors, part loading, etc. has numerous communication schemes for example, CANOpen, DeviceNet, FOUNDATION Fieldbus, Interbus-S, LonWorks, Profibus-DP, and SDS. Home automation—scheduling and automatic operation of water sprinkling, heating and air conditioning, window coverings, security systems, lighting, etc.—is being fought over by wireless solutions including WiFi, Zigbee, Z-Wave, and BlueTooth as well as wire solutions including HomePlug (over AC wiring) and HomePNA (over phone lines).

Massimini believes the unifying force bringing these disparate communications schemes together is the Internet Protocol version 6 (IPv6), the latest version of the Internet Protocol (IP). The communications protocol provides an identification and location system for computers on networks and routes traffic across the Internet. IPv6 was developed by the Internet Engineering Task Force (IETF) to deal with the long-anticipated problem of IPv4 address exhaustion. If Cisco's estimate of 25 billion devices connected to the Internet by 2015 and 50 billion by 2020, IPv6 is not a minute too soon. As to how these billions of IoT devices will communicate, Massimini cites the emergence of reference designs from OEMs including Qualcomm, Broadcom, TI, Freescale, ST, and others that provide intelligent gateways to bring all these devices together and provide IPv6 traffic to where ever on the network.

Once that problem is solved, Massimini sees an even greater one rearing its head: security. The cautionary tale he uses to illustrate this danger is the hacking attack that showed the gaping hole in retailer Target's network security. The attack originated from Fazio Mechanical Services (FMS), a Sharpsburg, PA-based heating, ventilation, and air conditioning (HVAC) systems that contracted to Target to provide not only HVAC installation and maintenance but also to monitor and control the environment with Target's retail outlets. The HVAC system can be accessed via an IP address. Somehow the hackers acquired the encryption key from FMS required to access Target network connecting point-of-sales (POS) terminals and were able to plant malware that copied every credit card transaction in the POS terminal where it was collected and transmitted the information to servers located at different locations around the globe.

According to the Symantic white paper, "A Special Report on Attacks on Point of Sales Systems" this is not an uncommon occurrence as the software to pull this off is readily available on the web and the incidence are not new as the first happened in 2005, when 170 million card numbers were stolen. Since the POS system cannot be network-segmented from other networks, Massimini says the solution that seems to be emerging is the replacement of magnetic strip credit and debit cards with smart cards like those used in Europe that employ the Europay, Mastercard and VISA (EMV) set of standards for card payments. EMV employs an embedded processor with strong transaction security features to protect card data.

Massimini says this lesson hasn't been lost on OEMs building intelligent IoT gateways and devices who are incorporating crypto engines of their own design in the microcontrollers controlling these products. This additional security may be late in coming as attacks are already beginning to occur in home according to Proofpoint, Inc. The security-as-a-service provider based in Sunnyvale CA claimed to have discovered the first proven Internet of Things (IoT)-based cyberattack. The company's press release reported 750,000 malicious email communications coming from more than 100,000 everyday consumer gadgets—home-networking routers, connected multi-media centers, televisions and at least one refrigerator—that had been used to launch attacks.

Implementing more layers of security in the end devices and the gateways they connect to will be costly. The commercial segments are most likely to accept this cost since there is an immediate benefit to the bottom line. Providing more security for consumer devices is problematic. The intelligent gateways for home will need to be the first line of defense. Keeping these security measures up to date will be another business service.

The Internet of Things is just the latest incarnation in the evolution of computers and communications. As the consumer demand grows for the benefits provided by smart connected devices, hardware and software vendors will build the affordable secure devices these consumers will buy.

Jonah McLeod is the Corporate Marketing Communications Director for Kilopass Technology Inc.

The Industrial Computer Afterlife

Legacy hard drives are dying and replacements are getting harder to find.

By Ben Hensley, Onyx Automation

Of all things in this world counted on to endure, steadfastly and reliably, hard drives are not among them. Like trolls they lurk, deep in the bowels of equipment, waiting for long weekends or holidays to spin themselves to pieces. I wince every time I switch on old equipment, knowing one day I'll lose the power-up Russian roulette and be left with an ominously, uselessly clicking disk. Even the soft, reassuring blanket of frequent backups is little comfort for those that maintain legacy hardware, because their disks are from another era: one of rattling ball bearings, deprecated interfaces and BIOS size limitations.

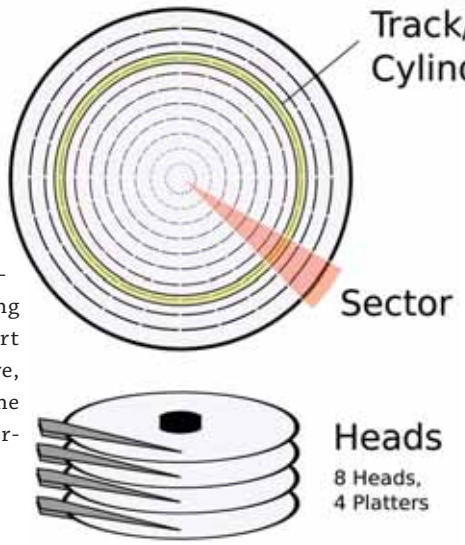

Figure 1. The logical and physical orientation of a rotating, mechanical HDD. (Courtesy: Wikimedia Commons.)

It Don't Come Easy

Replacing a failed legacy hard disk is nontrivial for two reasons: first, a majority of drives installed in industrial environments from the '80s until fairly recently use the IDE (now called parallel-ATA or PATA) interface. Parallel-ATA was the standard for desktop machines from the mid-1980s to its replacement by serial-ATA (SATA) in 2003. Though new parallel-ATA disks are still being manufactured, their production numbers are decreasing, and serial-ATA drives reached 99% of the desktop market share in 2008. OEMs cannot reasonably be expected to continue supplying this older technology indefinitely.

The second, and perhaps more infuriating reason it's difficult to replace an old hard disk, is modern drives have too much capacity: one does not simply put a 320GB hard disk into a 386-era machine. Depending upon age and manufacturer, a legacy computer may fail to recognize a disk drive that is larger than some seemingly arbitrary size. To understand these capacity limitations we need to take a look at the computer's BIOS—the on-board firmware that initializes and tests memory and peripherals before loading the operating system.

BIOS Bias

One of the functions a computer's BIOS supplies to older, real-mode operating systems like DOS is a hardware abstraction layer (HAL) to hide differences between configurations from the software. When a program needed to access the hard drive, rather than needing to know the exact specifications of the disk, it would just call interrupt (INT) 13h and start using the BIOS for disk reads and writes. This saved programmers from needing to write for any possible hardware configuration, as they could use the standard interface provided by the BIOS.

To the BIOS, the hard drive is a collection of blocks (they were 512 bytes, now 4 KB is standard) arranged around the surface of the disk platters (Figure 1). Each block has its position uniquely identified using the cylinder/head/sector (CHS) addressing scheme. CHS is possibly the only implementation of cylindrical coordinates I've seen since college, and, like most uses of cylindrical coordinates, probably seemed clever at the time but ended up being a bad move.

The cylinders are concentric rings from the spindle to the edge of the disk platter, and the sectors are the position around the disk. Heads are the actual read/write head, with one on the top and one on the bottom of each platter (Figure 2). By specifying a cylinder, head and sector (or r, h, and θ), any block could be accessed. Space being at a premium in the BIOS code, ten bits were used to store the cylinder number (0-1023), eight bits for the head (0-255) and six bits for the sector (zero is not used, so 1-63). These values, when combined, can address a maximum of 16,515,072 blocks. At 512 bytes each, the maximum usable size should be 8.4GB; however, other limitations would show up along the way.

In the mid-1980s, Western Digital developed the Integrated Drive Electronics (IDE) standard which moved the hard drive control electronics from the motherboard to the drive. This greatly simplified controller design—no need to support different types of drives—as well as motherboard design—no need to deal with the

Figure 2. Inside a mechanical disk drive. (Image courtesy of William Warby; http://bit.ly/1iaFTed.)

internals of the drive such as spinning the platters and moving the heads. All that was necessary in the new configuration was for the host to ask for a particular block and the drive would read or write to it as needed. As before, these blocks were addressed by a cylinder, head and sector, but the IDE specification used a different number of bits to store each number: 16 bits for the cylinder number (0-65535), four bits for the head (0-15) and eight bits for the sector (again, zero is not used, so 1-255). These values can address a maximum of 267,386,880 blocks, for a maximum of 137GB.

Unfortunately, for the operating system to talk to the BIOS and for the BIOS to talk to the disk, both limitations must be met, so the combined maximum is 1024 cylinders, 16 heads, and 63 sectors, addressing 1,032,192 blocks, or 528MB. This was the first common restriction encountered, but it wasn't a problem until around 1994 when disk drive capacity began to reach and exceed that size. One of the first workarounds was developed by Phoenix Technologies, a BIOS manufacturer: by adding a translation layer between the BIOS CHS parameters and the IDE CHS parameters, the full 8.4GB addressable by the BIOS could be used.

With the difficulty in finding replacements for motherboards or BIOS, and the cost associated with replacing and upgrading existing equipment, it's natural to focus on the hard disk as a cost-effective way of extending the life of embedded and industrial equipment.

BIOS Bit Shifting Scheme

The translation used bit shifting—a fast operation—to multiply or divide the physical cylinders and heads (as reported by the drive) to logical cylinders and heads (as used by the BIOS). For example, a disk that had 16,384 physical cylinders and 16 physical heads would have the cylinder number divided by 16 while the head number would be multiplied by 16, giving a logical geometry of 1024 cylinders and 256 heads. This allowed the operating system to use a logical CHS address that didn't exceed the maximum number of cylinders when talking to the BIOS, and the BIOS could in turn talk to the disk drive without exceeding the maximum number of heads.

There were, however, two other barriers that weren't completely fixed by Phoenix's translation scheme. Some BIOSes allocated only 12 bits for physical cylinders (0-4095), limiting those machines to 2.1GB disks. The second limitation wasn't actually a fault of the BIOS, but rather of DOS not supporting drives with 256 logical heads. DOS stores the number of drive heads as an eight-bit number (0-255); if the BIOS reports 256 heads, only the lower eight bits are used and DOS sees a disk with zero heads. As most drives larger than 4.2GB came through the translation layer with 256 heads, this was an effective limit for DOS-based operating systems until another rule was later added to the translation scheme to cover this specific case. Depending upon the BIOS manufacturer and release date, finding a drop-in replacement drive can be a peculiar mix of eBay, trial-and-error and hair-pulling.

The usual advice given for BIOS-related limitations is to upgrade: replace or reflash the BIOS, replace the motherboard or replace the entire machine. Finding an updated BIOS for a specific legacy motherboard, if one was ever offered, grows exponentially harder as time passes and may not be possible after ten to fifteen years. Even if the manufacturer is still in business, chances are the product reached the end of its support life long ago. Replacing a motherboard with a newer version that would support new, large-capacity disks might mean the loss of any on-board proprietary hardware, not to mention those once-ubiquitous ISA slots.

ISA was a data bus standard used for years to add to a computer expansion cards such as digital I/O, stepper motor control, video capture and more, but now ISA slots on consumer motherboards are unheard of. Some industrial computing companies do offer solutions to build a brand-new, x86-compatible machine that will allow old hardware to be connected to a modern BIOS that supports such niceties as large disks and USB ports. Unfortunately, bringing software from an era of megahertz clock speeds into the gigahertz world isn't always possible. Even with tools to slow current processors, issues of timing, especially with old expansion cards, can make porting to current hardware impossible.

The Legacy HDD Plan

Some disaster recovery plan needs to be made, though, because all hard disks will have to be replaced eventually. In large studies of component failures, like the 2007 Carnegie Mellon paper—Disk failures in the real world: What does an MTTF of 1,000,000 hours mean to you?—hard disks are always high, if not the highest, on the list of components most likely to fail. Disks manufactured before the early 2000s are at even more of a disadvantage, for in addition to their age, they still used ball bearings to spin the disk platters. New drives use fluid dynamic bearings due to the increasing vibration, loss of running accuracy, fatigue flaking of surfaces or any number of dire fates that await aging ball bearings.

"Whatever has a beginning must also end."
-Melissus of Samos
"Everything dies."
-Peter Steele

With the difficulty in finding replacements for motherboards or BIOS, and the cost associated with replacing and upgrading

existing equipment, it's natural to focus on the hard disk as a cost-effective way of extending the life of embedded and industrial equipment. Hard disks are low-hanging fruit: a high failure-rate, critical part that's designed to be easy and cheap to replace. But with drop-in replacements not available for most legacy equipment, how are we to proceed?

Solid-state disks (SSD)—those based not on writing to magnetic media but on storing electric charge—have been around in some form since at least the 1970s. Early devices were volatile, like RAM, and data was not retained if power was lost unless backed up by battery or other power supply. In the 1980s, flash-based storage was introduced—a non-volatile technology that did not need battery backup to retain its data. Though slower than RAM-based devices, their resilience to shock, vibration and temperature extremes made them popular replacements for mechanical hard drives in critical applications.

When solid-state disks started to become widely available, their low density and small capacities made them a seemingly perfect answer for replacing disks in machines with BIOS capacity limitations, as one could easily find a 128MB or 256MB PATA SSD. The drive controller for a solid-state disk operates much differently than one from a mechanical disk, but they both present the same interface to the BIOS, and thus are virtually indistinguishable. Unfortunately, the smallest solid-state disks found in the market today are already 4GB or larger, and Moore's Law will continue to march us onward toward more transistors, smaller sizes and further out of the reach of our legacy equipment.

To address issues of hard disk replacement at Onyx Automation, we use custom-sized solid-state disks. With high-reliability, single-level cell (SLC) flash technology, we build and supply disks as small as 128MB with no moving parts and smaller power requirements than mechanical hard drives. Our services include drive imaging for disaster recovery, and migration of existing data to upgraded equipment. Onyx Automation has over two decades of experience in moving old operating systems to newer hardware, including MS-DOS, Japanese DOS, Japanese OS/2, Windows 3.11/NT/XP, and Linux.

Ben Hensley, owner of Onyx Automation, has been performing equipment repair and upgrades for the semiconductor industry since 2006, combining new hardware with old operating systems. He has a degree in chemical and biomolecular engineering from the Georgia Institute of Technology. Ben can be reached at: ben@onyxautomation.com

MH70I Modular Industrial PC

Compatible Operating Systems: Windows Embedded and Linux

Supported Architectures: MEN Micro's F22P - 3U CompactPCI® PlusIO Intel® Core™ i7 CPU Board (Intel® Core™ i7-3517UE or Intel® Celeron® 1047UE), 16-bit, Embedded Intel® (Pentium, Embedded Intel® Architecture etc.)

Supported Bus Structure: PICMG

MEN Micro's MH70I modular application-ready industrial PC is configurable with a wide selection of standard hardware, software and accessories to provide cost-effective customization.

The MH70I owes its modularity to the combination of cards that make up the heart of the unit. In addition to the system slot, the MH70I includes two each of CompactPCI and CompactPCI Serial peripheral slots and two each of PCI and PCI Express slots, totaling nine individually-configurable slots in the system. Two additional PSU slots supporting AC, DC and UPS, ensure redundant, reliable operation.

The unit's compact half 19" rack size allows two systems to be placed on one rack side by side. The MH70I can be either rack- or wall-mounted with natural convection cooling or with an added fan tray at the bottom of the system, depending on the mounting position.

The standard configuration comes with MEN Micro's F22P CompactPCI PlusIO SBC using an Intel Core i7 or Celeron processor with one VGA Interface, two USB ports and two Gigabit Ethernet interfaces on the front. It also includes up to 16 GB of DDR3 DRAM with ECC.

Various peripheral boards are available, including analog or binary I/O via M-Modules, fieldbus functions, SATA hard disks or wireless functions as well as Ethernet interface boards or Ethernet switches. A SATA RAID with up to four HDD shuttles can be built using the CompactPCI Serial slots.

FEATURES & BENEFITS

- Compact 19" application-ready system
- Rack-mounted or wall-mounted
- 2 CompactPCI® slots for fieldbus functions, RS232, analog I/O, digital I/O, Ethernet
- 2 CompactPCI® Serial slots for SATA RAIDs, Ethernet
- 2 PCI or PCI Express® slots for half-length cards

TECHNICAL SPECS

- Up to nine configurable slots in one system
- Half 19" Rack Size with Cooling By Natural Convection or Fan Tray
- Operating temperature: 0°C to +50°C
- Up to 16 GB soldered DDR3 DRAM with ECC
- Up to four SATA hard disks for RAID

Vertical Market Applications: Industrial, Transportation

CONTACT INFORMATION

MEN Micro Inc
860 Penllyn Blue Bell Pike
Blue Bell, PA 19422
United States
tel: 215-542-9575
fax: 215-542-9577
sales@menmicro.com
www.menmicro.com

Two Premier Conferences Showcasing the Embedded Systems Industry

10th ANNUAL
MULTICORE
DEVELOPERS CONFERENCE

Resolving the Technical Aspects and Business Challenges of Designing with Multicore Processors

IoT
INTERNET OF THINGS
DEVELOPERS CONFERENCE

Resolving the Technical and Business Challenges of Getting Connected to the Internet of Things

Plan now to attend! MAY 6 -7, 2015 Santa Clara, CA USA

For information on exhibiting or sponsoring contact:
Clair Bright +1 415-225-0390 x15 or cbright@extensionmedia.com