HIGH-PERFORMANCE EDGE COMPUTING

The Frontier for Supercomputing and Acceleration

Exclusive License to Distribute: Eurotech
By Dan Mandell, Senior Analyst
and Chris Rommel, Executive Vice President

THE VIEW INSIDE

Edge computing is the new domain for innovation. Next-generation analytics, machine learning (ML), and other high-performance workload processing require comprehensive intelligent edge frameworks and platforms. Furthermore, these platforms must have maintained pace with the convergence of embedded and enterprise/IT domains through the past several years. High-performance embedded computing (HPEC) features distinct development requirements and challenges spanning hardware, software, connectivity, platform integrations, and security. Years ago, a few suppliers had the foresight and ambition to begin tackling high-performance edge computing and forming a partnership ecosystem including key software and technology providers. Those that did, however, are now reaping the fruits of their labor and have a dramatic head start over competitors scrambling to take a stake in a critical piece of the market for many different industries.

EDGE COMPUTING TAKES CENTER STAGE

Industries are Now Data Dependent

Data is a critical resource. It fuels emerging high-value applications across many industries. This business lifeblood is being generated at an increasing rate thanks to the growing availability and capabilities of industrial sensors, smart cameras, imagers, and other data inputs. The evolution of embedded vision and sensing technology is sweeping many industries, enabling new possibilities with regard to remote monitoring, data analytics, ML, and real-time control. The enhanced features of (increasingly lower-cost) connected sensors and cameras is fueling the development of low-latency architectures, more broadly pushing the bounds of deterministic processing functionality beyond mission-critical workloads.

Exhibit 1: Leading Industry Edge Applications

Industry	High-Performance Edge Applications
Aerospace & Defense	C2/C4 On the Move & C4ISR Situational Awareness, Electronic Warfare, Ground Based/Mobile & Airborne Radar Processor, Navtronics, UAVs
Automotive	Autonomous Driving, Intelligent Mobility, Infotainment, Real-time Mapping, Sensor Fusion (radar/LIDAR, vision, etc.)
Energy & Utilities	Energy Management & Response, Oil & Gas Exploration, Power Plants, Smart Grid & Power Distribution
Industrial Automation & Control	Autonomous Robotics, Harsh/Dusty Environment Data-Intensive/ Imaging Processing, Real-time Equipment Monitoring & Diagnostics
Medical & Healthcare	Connected Healthcare, Notifications, Patient Data Management, Predictive Analytics & Control, Remote Device Management
Rail & Transportation	Passenger Counting, Railway Management, Rolling Stock
Retail Automation	AR/VR, Video Surveillance & People Counting/Analytics
Other	Agriculture, Cyber-Security & Networking, Geological Surveys & Exploration, Weather & Environmental Condition Monitoring

A major challenge still facing embedded technology suppliers and manufacturers is the end-to-end integration of connected systems, cloud endpoints, and third-party platforms or services. A flexible embedded framework is required to ensure maximum use of data generated/collected over the long term, particularly for fueling business applications and cross-integration with different IoT platforms or PaaS/SaaS. This is a necessary precursor to developing high-performance edge servers and processing solutions. The framework must thread through the OT and IT domains to include support for various field infrastructure, communications infrastructure, (IoT) application infrastructure, as well as the application/enterprise IT layer. Without a flexible device-to-cloud integration platform and supporting software/middleware libraries for gateways and other connected systems, solutions for the high-performance edge cannot scale or adapt to the dynamic requirements of solution providers and end users.

High-Performance Edge is Redefining Boundaries

The network edge is the current incubator for platform innovation. While the specific definition and bounds of the "edge" are up for debate between leading technology providers and standardization bodies, its growing importance in industry solution architectures is undeniable. Through the past several years, a growing number of device classes have supported the evolving edge domain, including traditional smart routers, M2M gateways, intelligent gateways, edge servers, and more. Each has played a significant role in enabling advances in data connectivity, management, and security. Growing data dependency is pushing requirements not only for data access, but also for throughput and processing at the edge, requiring robust computing architectures.

IoT routers and M2M gateways were the first on the edge scene more than a decade ago to tackle the challenges of heterogeneous connectivity and data management. As the need for greater intelligence and flexibility from the network edge grew, intelligent gateways were deployed with new capabilities for API integrations and embedded application processing/analytics. More recently, field-deployable edge servers and hyperconverged acceleration platforms have emerged to bring more traditional datacenter technologies closer to the operational side of deployments to securely partition, process, and manage various "heavy" workloads or large densities of connected OT infrastructure. For current and future deployments, the spotlight is on the quickly growing market for intelligent gateways, edge servers, and other high-performance edge hardware platforms to support next-generation applications.

Exhibit 2: Global Shipments of Fixed Edge Servers & Edge Hardware Platforms, 2017 & 2022
(Billions of Dollars)

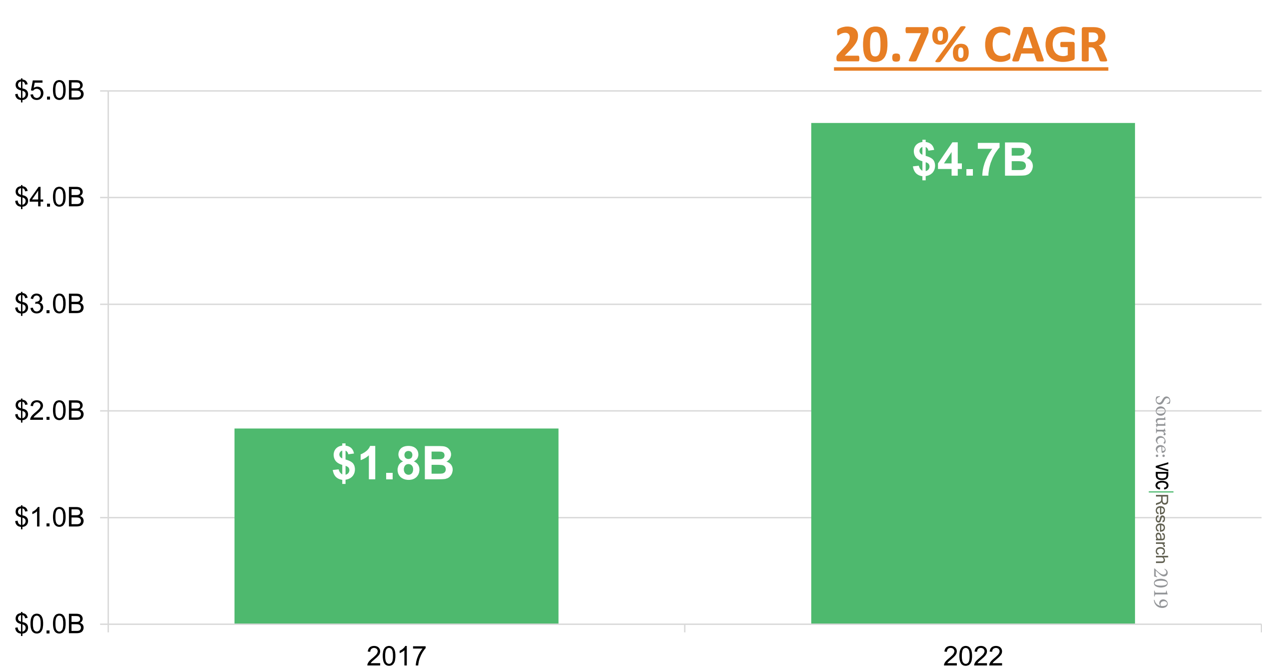

High-performance edge infrastructure is drastically changing device-to-cloud solution architectures in the development of new use cases leveraging connected data streams. New classes of hardware, software, and IP are needed for hosting increasingly popular accelerated processing, machine learning/inference, and vision applications at the edge. In some instances, these growing requirements are fueling the demand for more vertically tailored features, functionality, computing profiles, and form factors. The intelligent edge requires a healthy blend of OT and IT infrastructure to stay at pace with system and solution requirements.

Hybridization of Embedded & Datacenter Technologies

The intelligent edge is the setting for the widespread convergence of the embedded and OT worlds with the enterprise/IT domain. This integration has driven the formation of countless new partnerships, collaborations, industry/technology groups and alliances, and standards for organizations of all sizes. Even those organizations with a strong history within the enterprise/IT domain such as Cisco, Dell, and IBM have had to expand their strategic collaborations and support given the incredible diversity and fragmentation of solution requirements (even among industry applications).

The majority of OT engineering teams and IT groups are wading into new territory working with each other—let alone with other third-party organizations. New integrations linking both worlds are needed to facilitate OT infrastructure management and edge application enablement. Linking both worlds features incredible challenges regarding security and the general inertia of OT and IT teams working by themselves for so long. For instance, high-performance edge platforms often require fast storage (e.g., with NVMe management) and highly configurable I/O cards. Not many technology providers from the embedded domain understand the complexity of deploying multiple device clouds and enabling managed services.

The Industry Needs Guidance & Direction

Very few organizations have the history and experience of integrating embedded/OT and IT platforms themselves to provide a flexible end-to-end solution architecture for development. Many embedded hardware suppliers have opted to stick to their traditional focus of developing new products and relying on systems integrators and/or a general "ecosystem" for end-to-end solutions development. However, developing solutions for the high-performance edge features a number of unique challenges that demand an experienced partner from both the OT and IT worlds. Requirements for deploying next-generation analytics and applications span traditional SWaP, thermal management, ruggedization, hardware acceleration, software development, security, and much more. OEMs and other enterprises looking to develop HPEC solutions are wading into new territory and need as many building blocks and as much professional support as possible.

Exhibit 3: Eurotech HPEC System

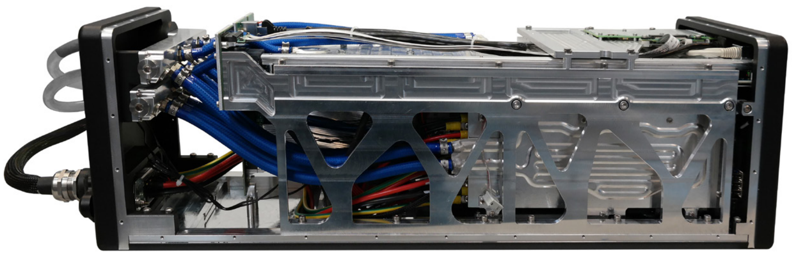

These HPEC building blocks (including hardware, software, and cloud) must be modular with the flexibility to support a wide range of configurations to suit a variety of developer preferences, workflows, and stack requirements. Simply supporting a list of different software frameworks, hardware platforms, and vendor plug-ins is not sufficient for most engineering organizations. The resources available must be easily combined/integrated with each other to reduce time-to-market and development costs for the target applications. Real, tangible benefits are often most easily gained through a homogeneous development platform and environment.

HIGH-PERFORMANCE EDGE FUELS UNIQUE REQUIREMENTS

Power Consumption, Ruggedization, & Thermal Management Hurdles

The high-performance edge features a unique blend of hardware requirements, with the biggest challenges stemming from adapting to environments beyond traditional server closets and controlled environments. Minimizing size, weight, power, and cost matter more for demanding HPEC applications with space and energy constraints than they do for enterprise/IT deployments. Even those few organizations that have played in high-performance computing and supercomputing have, for the most part, not provided enough emphasis on minimizing power consumption or ruggedizing their hardware, which are imperative for HPEC at the edge. Edge infrastructure effectively comprises OT systems often requiring similar protections as downstream devices and sensors from harsh environments against dust/water ingress, vibration, EMI, temperature, etc. In some industries such as automotive and industrial automation, certification with industry standards adds another layer of requirements and challenges for HPEC.

The edge also features unique challenges regarding cooling, as the system itself must maintain thermal management using passive, water, or other configurations in the field to sustain the need for speed. In fact, many of the emerging high-value HPEC market opportunities are popping up in environments where IP-rated/sealed enclosures are required to protect embedded systems, further complicating power and cooling design requirements. Accruing valuable systems IP for (high-performance) ruggedized cooling, particularly in mobile environments, requires years of experience. For example, companies like Eurotech have broad experience with embedded liquid cooling mechanisms for its HPEC hardware in autonomous vehicle applications that can be connected to vehicle cooling systems.

HPEC Software Requirements Differ from IoT

Edge technology providers are looking to bring more datacenter performance and functionality closer to the field to enable new possibilities for application, data, and device management. For enterprise-class hardware at the edge, server-type technology and solutions are needed. Flexible hardware architectures are essential with intelligent docks and hot-swappable components to adapt to a variety of deployment types or workloads. Virtualization can be a powerful solution for both embedded engineering and enterprise/IT systems. Virtualization provides a critical abstraction layer between hardware and software enabling the deployment of new applications on legacy platforms. Virtualization support is of growing importance and value to many industries, particularly in safety-critical industries such as aerospace and defense, automotive/rail, industrial automation, and medical.

Exhibit 4: Current and Expected Future use of Hypervisor/Virtualization/Containerization (Percentage of Respondents)

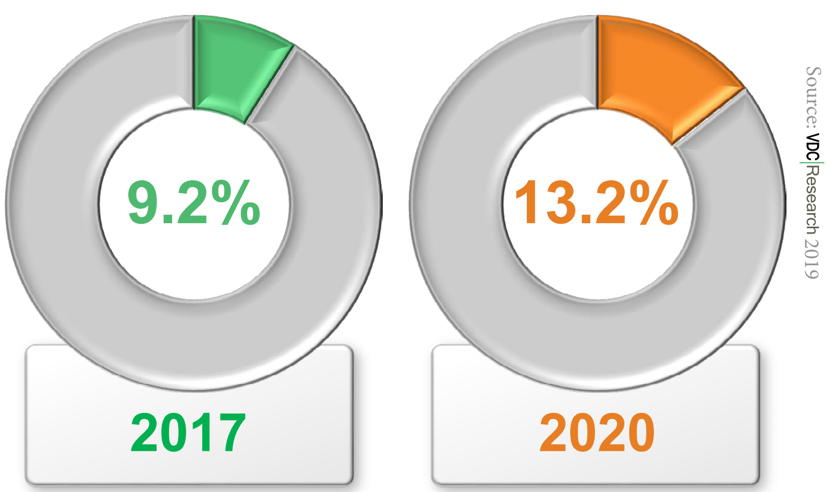

A comprehensive development framework is critical for high-performance edge systems to ensure greater software flexibility and control for device (lifecycle) management. The framework must branch its support in many directions to maximize its deployable footprint supporting different programming languages, cloud environments, field equipment, and IoT platform integrations. Developers require the connectivity and interface library support with leading field devices and protocols for their target applications and geographies. Further, embedded frameworks need to support high-level languages such as Java to maintain compatibility with shifting programming language (and API) use and to leverage a larger developer ecosystem. OSGi-based frameworks in particular have generated considerable traction in recent years enabling the development and deployment of software modules while leveraging the rich Java developer ecosystem.

Exhibit 5: Programming Languages Used in Current Projects
(Percentage of Respondents, Top 10 Selections Shown)

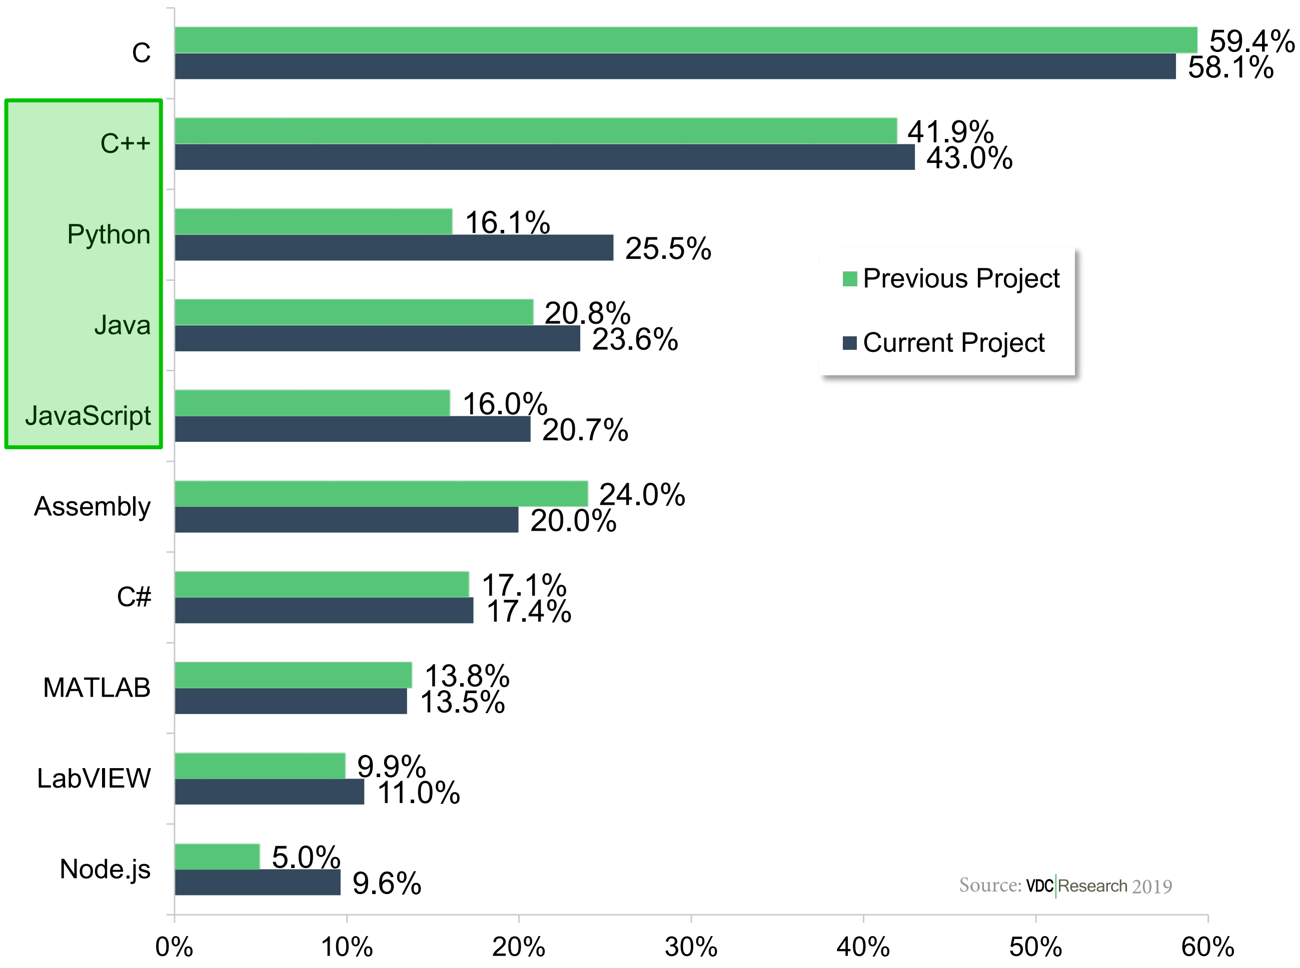

Acceleration is a Must

The demand for highly parallel processing capabilities using hardware acceleration is growing across different industries to tackle a variety of heavy workloads ranging from AI, to cryptography and machine vision, among others. A variety of processor types spanning CPU co-processors, FPGA, and GPGPUs are enabling hardware acceleration. Each technology has its own advantages and drawbacks with regard to their use for hardware acceleration. There are also key considerations to be made regarding the changing dynamics of the development/deployment of accelerators integrated on SoCs, board-level solutions, and integrated systems mixing different cores and IP for processing application data.

Embedded GPGPUs have emerged as a frontrunner for hardware acceleration because of the inherent performance of the underlying architecture as well as strong support from the two suppliers with 95%+ share in the GPU market—AMD and NVIDIA. Leaders in the single-board computer and blades market such as Abaco Systems, Artesyn Embedded Technologies, and Eurotech have each launched offerings that leverage GPGPU technology to tackle mission-critical workloads and applications. In fact, 5G & High-Performance Edge Infrastructure featuring GPGPUs are growing at a CAGR of 15.0% from 2017 through 2022—faster than shipments with embedded FPGAs. In June 2018, for the first time in history, most of the flops added to the TOP 500 rankings came from GPGPUs and not CPUs. In fact, in the latest rankings published in November 2018, three of the top five ranked systems feature GPGPU accelerators. Given the rich ecosystem of core GPU technology and tools providers with the growing support from hardware suppliers featuring different levels of off-the-shelf integration and performance, we expect GPGPUs will continue to be an accelerator of choice for many different industries.

BUILDING BLOCKS & DEVELOPMENT FOR HPEC

Rich Developer Support Needed for HPEC Processors, Boards, & Systems

HPEC development is a multi-tiered challenge, requiring strong support from a number of organizations to be effective. First are the embedded processor and IP providers like Arm and Intel, which are steadily extending their support and influence of system stacks and solutions. They are looking to offer supporting components for cloud connectors, connectivity/interfaces, interconnects and networking, machine vision (e.g., Movidius), on-chip analytics, security, server management, SoC core management (e.g., hybrid architectures with CPU+FPGA, CPU+GPU), and virtualization, among others. Growing real-time processing requirements are also affecting the selection processes for embedded processors, OS, and middleware to enable better decision-making and high-value analytics.

As discussed, thermal management is a major hurdle that generally raises with higher-performing processors for edge servers and other hardware. Systems engineering for the datacenter is very different from the edge environment. Furthermore, cooling at the edge is a brand new greenfield for traditional enterprise/IT infrastructure providers. For example, only through liquid cooling, which is a tremendous challenge to enable in an embedded/sealed environment while meeting reliability and ruggedization (shock, vibration, etc.) requirements, can many deployments offer the computation density required for the latest innovative AI applications. For autonomous driving, liquid cooling enables extended temperature ranges in both directions to warm or cool hardware (server) components in embedded fields. This new class of hardware requires innovative power management solutions to balance the power-hungry CPU/APU and GPGPU hardware of the datacenter with the power and certification requirements of field deployments.

Other higher-level building blocks are available from boards and systems providers to facilitate edge application enablement of end-to-end or edge-to-cloud solution architectures. High-end building blocks for HPEC range from a variety of components such as OS support and software optimizations, containers and orchestration middleware, to system management solutions, functional safety tools, and frameworks for enabling hybrid edge clouds. The building blocks will only be effective, though, within an open, modular, and scalable hardware and software platform. The need for a versatile platform has fueled the growth of open-source projects such as Eclipse Kura, which offers a bundle of features and API access to the hardware interfaces of IoT gateways and other edge infrastructure. High-speed interconnect standards like CCIX for hardware acceleration are also generating traction as it becomes a more important element of edge designs. Efficient HPEC development requires coordination with a number of different entities (and their roadmaps).

Nobody Can Do it Alone

OEMs need to effectively "choose their spots" for in-house development after deciding which building blocks they want to use to jumpstart development. Failure to do so will inevitably lead to project delays with the most frequently cited difficulties being "Technical Obstacles" (24.8% of respondents), "Changes in Specifications" (23.0%), and "Complexity of the Application/Technology" (22.7%).[1] In the same light, embedded hardware providers must ensure they have enough building blocks either internally developed or available through partnerships and third-party platform/software support. These alignments are often strategic collaborations between companies such as IBM, VMware, Microsoft, Oracle, or Red Hat, bringing together more design elements for end-to-end solutions while enabling more groups of software developers and engineers to work together. The proliferation and growing capabilities of embedded Linux and its supporting ecosystem throughout high-performance edge infrastructure are another example of a foundational element for development upon which most suppliers and end users are looking to build for more industrial use cases.

> "*81% of 5G & High-Performance Edge Infrastructure hardware revenue shipments in 2018 used Linux-based operating systems.* – VDC Research"

Security Cannot be Overlooked

Security is now at the forefront of design planning and vendor evaluations. OEMs and systems integrators need wide coverage of embedded hardware and software security and data protection throughout the deployed lifetimes of their product systems. For the intelligent edge, embedded security requirements are most prominent because of the expanded attack surface exposed by connected systems and infrastructure. The need for security stretches throughout the solution architecture to validation and authentication of identities for devices and platforms, integrated PKI and certificate management, encrypted communications, secure pairing, and the provisioning of devices. A secure software management and distribution platform is essential to maintaining security health throughout the deployment. Fortunately, as edge systems adopt more datacenter technologies, they are also taking on more security functionality such as network and system security (firewall, hardening), role-based access control, secure management access, virtual machines, and VPN services for remote management.

Exhibit 6: Stack Component Requirements (Aside from OS) Featuring the Greatest Expected Growth in Use from Current Projects to Three Years From Now
(Percentage of Respondents)

Stack Components	Growth in Use
Embedded Security Hardware (e.g. Hardware root-of-trust, cryptographic coprocessor, biometric sensor, etc.)	**+14.1%**
Cloud Agent	**+12.4%**
Embedded Security Software (e.g., Anti-virus, cryptography, authentication, etc.)	**+10.9%**
Firmware Over-the-Air (FOTA) Software	**+9.0%**

[1] *2019 Embedded Boards, Modules, and Systems Report* (VDC Research)

VDC'S VIEW

The Edge Features Many Evolving Requirements

HPEC for the edge is a tumultuous landscape currently with a range of dynamic requirements for hardware, software/middleware, connectivity, and cloud support to fuel next generation functionality and applications. The migration of datacenter technologies and hyperconverged systems to the edge is radically changing the status quo for high performance embedded platforms. For hardware, demand is ramping for fully validated and optimized GPGPU solutions, with varying acceleration and scalar performance from all types of suppliers from semiconductors to integrated/OEM edge servers. However, developers need the most support where developer costs are growing fastest – for software, data/analytics, and cloud/IoT. OEMs, integrators, and others are looking more frequently for their embedded technology suppliers to help fill in these gaps in development expertise.

Exhibit 7: Estimated Distribution of Development Costs
(Average of Respondents)

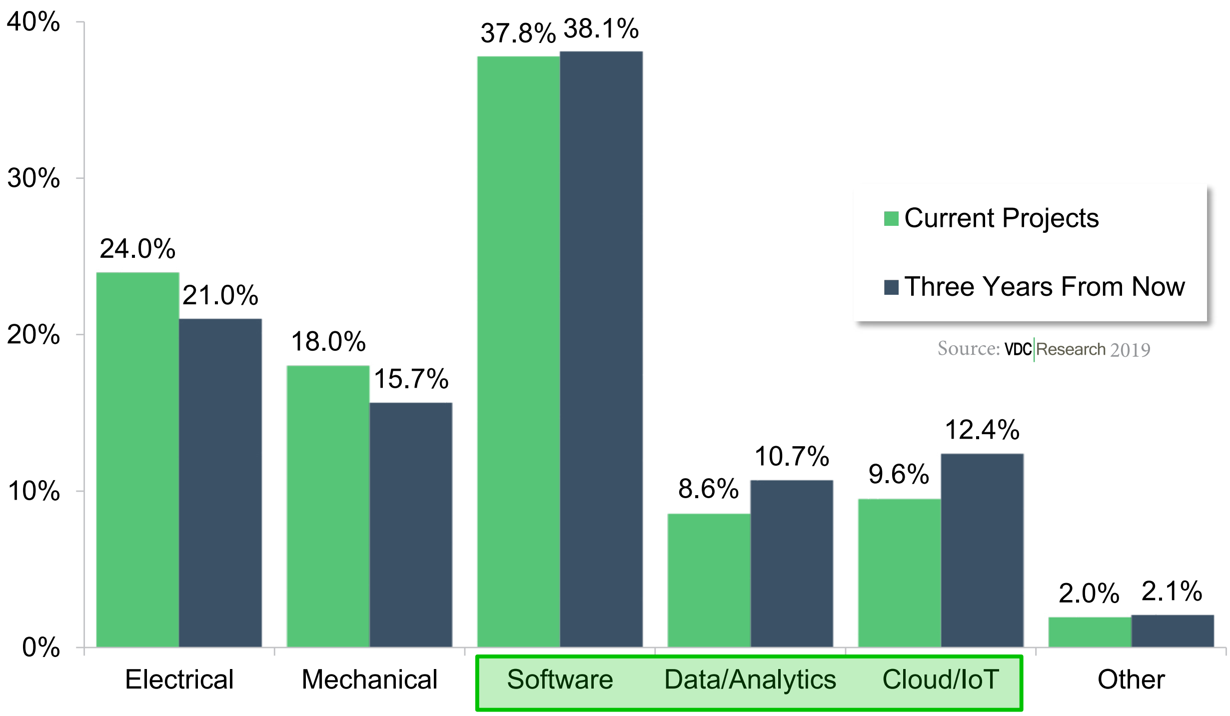

Nevertheless, more-traditional requirements for hardware and I/O remain and evolve with the demands of HPEC. For example, high-speed PCIe Gen 4 lane bus technology features reduced system latency, double the transfer rate and throughput of Gen 3, and is able to achieve the internal performance required for NVMe storage and GPU computing applications (as long as signal integrity is preserved). Other technologies, such as real-time kinematic (RTK) positioning for autonomous driving, geological surveys, and other precise location-based applications, are also growing in value to support leading HPEC deployments at the edge.

Broadening Demand for Embedded Server-Class Technology & Development Support

HPEC is growing rapidly in a variety of domains for enabling applications ranging from AI/ML at the edge, autonomous driving, digital security and surveillance, to distributed radar, exploration, and smart agriculture. The vast amount of data being generated in the field, combined with increasingly capable off-the-shelf embedded platforms and frameworks, is enabling new high-performance workload processing in the field. The business value of enabling the intelligent edge is critical today across many industries for improving operational efficiency and establishing a competitive edge. For example, geological surveys—whether being done in the arctic or your home neighborhood—rely on enormous datasets that traditionally needed to be processed in a datacenter to provide insights for future measurements or action. With HPEC in the field, the surveyor can determine whether additional measurements should be made nearby based on real-time edge analytics and save the organization from a potentially costly follow-up deployment later on.

The high-performance edge demands the production of quasi OT/IT systems with requirements and features stemming from both domains. Fortunately, best practices and technologies from the datacenter domain are being brought into the embedded world from vendors such as Eurotech, which have straddled the two domains for many years. Intelligent system/server management and virtualization are critical to HPEC at the edge. Enabling sophisticated edge computing platforms also requires a cohesive partnership ecosystem across different solution/stack elements. The most versatile and future-proof edge platforms are those built on open standards and technologies. As with other types of embedded hardware, developer kits and reference hardware will be critical to mitigating the difficulties and challenges of adopting new technologies and features for edge developers. Further, ruggedization is a key element of HPEC and will be a prominent differentiator for computing hardware and components suppliers.

The Future – Edge Supercomputing & Autonomy

No matter how you look at it, the next generation of autonomous systems and distributed computing requires greater edge processing performance to evolve. Developing high-performance edge platforms, though, is much more challenging than what most engineers and developers have faced in the past with a hybridization of OT and IT requirements on shared equipment. At the same time, suppliers are tasked with providing increasing support for vertical/industry-specific expertise and functionality as projects become more complex. Few suppliers have the in-house experience, products, and support to enable the development of high-performance edge system platforms or solutions.

Suppliers with both HPEC and IT-related expertise stretching the bounds of embedded and enterprise technology are becoming increasingly valuable for edge supercomputing and autonomy. Edge servers and platforms demand the challenging combination of low power consumption, intelligent connected hardware, a flexible developer framework, a rich partnership ecosystem, open standards, and professional services for integrating user-selected accelerators and modules. To reap the benefits of the converging embedded/OT and enterprise/IT domains requires the harmonization of several different and new technologies from both domains using an open and scalable approach. The edge may represent the conceptual middle ground between end devices in the field and the datacenter/cloud, but its prominence for enabling the next wave of high-value applications and services is unmatched. The next big battleground for supercomputing and acceleration is at the high-performance edge.

ABOUT THE AUTHORS

Dan Mandell

Dan Mandell supports a variety of syndicated market research programs and custom consulting engagements in the IoT and Embedded Technology practice. He leads VDC's annual research services for embedded processors, boards, integrated systems, IoT gateways, and other computing hardware. Dan's insights help leading technology providers align their go-to-market planning and competitive strategies with the dynamic embedded landscape and its constantly evolving buyer behaviors, technology adoption, and application requirements. His working relationship with VDC dates back to 2005 and includes time supporting Business Development as well as the AutoID practice. Dan holds a B.S. in Information Systems Management from Bridgewater State University.

Contact Dan:

dmandell@vdcresearch.com

Chris Rommel

Chris Rommel is responsible for syndicated research and consulting engagements focused on development and deployment solutions for intelligent systems. He has helped a wide variety of clients respond to and capitalize on the leading trends impacting next-generation device markets, such as security, the Internet of Things, and M2M connectivity, as well as the growing need for system-level lifecycle management solutions. Chris has also led a range of proprietary consulting projects, including competitive analyses, strategic marketing initiative support, ecosystem development strategies, and vertical market opportunity assessments. Chris holds a B.A. in Business Economics and a B.A. in Public and Private Sector Organization from Brown University.

Contact Chris:

crommel@vdcresearch.com

ABOUT VDC RESEARCH

VDC Research
Insights for the Connected World

Founded in 1971, VDC Research provides in-depth insights to technology vendors, end users, and investors across the globe. As a market research and consulting firm, VDC's coverage of AutoID, enterprise mobility, industrial automation, and IoT and embedded technologies is among the most advanced in the industry, helping our clients make critical decisions with confidence. Offering syndicated reports and custom consultation, our methodologies consistently provide accurate forecasts and unmatched thought leadership for deeply technical markets. Located in Natick, Massachusetts, VDC prides itself on its close personal relationships with clients, delivering an attention to detail and a unique perspective that is second to none.