RadiSys
THE POWER OF WE

WHEN YOU'RE FINISHED READING
GO TO: http://www.radisys.com/wp
TO LET US KNOW WHAT YOU THOUGHT

WHITEPAPER

STANDARDS-BASED SYSTEMS FOR HIGH-END IMAGING APPLICATIONS

VERSION 1.0 | OCTOBER 2006

CONTENTS

- STANDARDS-BASED SYSTEMS FOR HIGH-END IMAGING APPLICATIONS 2
- A WIDE SPECTRUM OF APPLCIATIONS 2
- CHALLENGES FACED BY IMAGING SYSTEMS MANUFACTURERS 3
- STANDARDS-BASED SYSTEMS 3
- COMPUTE COMPLEXITY OF IMAGING APPS 4
- MULTI-PROCESSING APPROACHES 4
- HARDWARE ACCELERATORS 5
- ATX AND SSI..6
- SOLUTIONS PROVIDED BY RADISYS 7

STANDARDS-BASED SYSTEMS FOR HIGH-END IMAGING APPLICATIONS

High-end image processing has long been the domain of specialized and proprietary hardware. While this approach has delivered the required performance, it comes with some important drawbacks. First is that development of proprietary hardware is expensive and time-consuming. Second is that the software (or firmware) is generally not portable and with every change of hardware platform the software needs to be significantly recoded and re-tested. To a certain extent, the use of proprietary systems has been unavoidable, given the complexity of the required computations and practical limits on processing time. However, as standards-based systems continue to increase in performance, more and more image processing applications become possible on these platforms. By augmenting general-purpose systems with application-specific accelerators it is possible to design very versatile and economical image processing solutions. The advantages of standards-based systems over proprietary ones are significant. Software investments can be preserved from one generation to the next. System designers can tap into a large pool of available hardware design and software programming talent. Finally, component parts are generally available at much lower cost. The end result is that products reach the market faster and at lower cost.

A WIDE SPECTRUM OF APPLICATIONS

High-end digital image processing is at the core of many of today's sophisticated imaging products. These applications can be classified in two categories:

- Visualization
- Machine vision

Applications in the first category produce an image or set of images that are a product of filtering a set of original images according to a particular algorithm or processing a particular data set for visual representation.

Applications in the second category produce a decision based on a set of images using a specific algorithm to filter the image first, then extract certain parameters from it, and finally, produce a decision, using certain algorithms or artificial intelligence techniques, which interpret the extracted parameters. Hence, machine vision takes visualization one step further, eliminating the human element by automating interpretation.

Table 1 provides a listing of some prominent digital imaging applications in both categories.
The medical imaging industry, with applications predominantly in the visualization category, is seeing some

VISUALIZATION	MACHINE VISION
Diagnostic Ultrasound	Semiconductor Wafer Inspection
MRI (Magnetic Resonance Imaging)	Parts Inspection
CT (Computed Tomography)	Food Inspection
Diagnostic Digital X-Ray	Robotic Assembly
Scanning Electron Microscopy	Autonomous Vehicle Navigation
Various Military Imaging Applications	Fingerprint & Iris Scanning
Luggage, Cargo and Passenger Screening	Facial Recognition

Table 1. Prominent visualization and machine vision applications

of the most exciting developments in its history with 4D (Four Dimensional) ultrasound, making it possible not only to see a fetus or a heart in 3 dimensions, but also in motion. Other "modalities", such as MRI, CT and Digital X-Ray are also moving to 3D and, in some cases, 4D images. Also, multiple modalities are being combined to provide optimal images of multiple types of tissue. As one can imagine, the additional image complexity adds greatly to the computational requirements.

The security industry, with applications also predominantly in the visualization category, is likewise seeing an expansion in the use of high-end digital imaging as a way to prevent catastrophic events. In these applications, key drivers going forward are increased sensitivity to detect small and disguised contraband and also higher throughput to avoid bottlenecks at ports.

The visualization applications listed in Table 1 involve different digital signal processing tasks to be performed, ranging from noise reduction to solving complex inverse problems. Image processing techniques typically involve vector processing algorithms, such as correlation, convolution and FFT (Fast Fourier Transformation). For example, in cardiovascular X-ray systems, X-ray radiation is used to generate a video image of the heart and surrounding blood vessels. This allows a surgeon to "see what they are doing" and position a stent or other device exactly at a point of restricted blood flow. In order to protect the patient, and the surgical staff, X-ray radiation levels have to be kept to an absolute minimum level. As a result, the raw image data is extremely noisy. Advanced noise reduction algorithms are required to quickly process the raw data and deliver clear images, in near real-time, to a high resolution monitor at rates of around 30 frames/second.

CHALLENGES FACED BY IMAGING SYSTEM MANUFACTURERS

Manufacturers of high-end imaging solutions face many challenges on different fronts:

- They must design high-performance systems, which have the bandwidth to rapidly process large datasets and complex mathematical operations inherent to the applications.
- They are required to reduce time to market, or more importantly reduce time to revenue.
- They must be able to preserve their investments in application software and firmware when shifting to the next hardware platform in order to avoid lengthy recoding and testing.
- They must keep development costs low to maximize profits and reduce the time to recoup product development costs.
- They must be able to deliver solutions that can easily scale in performance, through different product configurations, field upgrades and by moving from one platform generation to the next.
- They must drive down system costs to be successful in competitive markets. In several segments of the medical imaging market for example, average selling prices "erode" by 5-10% per year, even as equipment performance improves.

Designing with proprietary systems is one way that systems manufacturers can meet their performance requirements. As mentioned above, this was historically the only viable approach for high speed systems. In addition to processing hardware, systems designers often resorted to proprietary interfaces to handle the high-speed I/O (input/output) requirements inherent in imaging systems.

However, the proprietary approach is one of diminishing returns and comes at a tremendous cost in the long term. Proprietary hardware constantly faces the risk of falling behind in the market and ultimately becoming obsolete. Application software and firmware portability tend to be severely limited and consequently time-to-market and time-to-revenue greatly suffer. Programming expertise on proprietary systems falls outside mainstream hardware and programming platforms and consequently the pool of knowledgeable hardware designers and software programmers is limited. Finally, proprietary systems also come with a higher price tag, since they almost always use lower volume parts than standards-based systems.

STANDARDS-BASED SYSTEMS

Historically mainstream x86 computer platforms could not meet the computational and I/O bandwidth challenges of high-end imaging applications. Legacy CPU architectures in combination with the interfaces such as ISA, PCI, PCI-X and ATAPI lacked both the computational horsepower and high-speed data transfer capabilities to process the large datasets involved in imaging within a reasonable amount of time.

Over time, mainstream systems have benefited greatly from "Moore's Law," which has resulted in processing power growing exponentially, almost like compound interest. Users of standards-based systems get these performance enhancements essentially for free, while those building proprietary systems generally have to make a significant investment to achieve each step of performance improvement.

A recent innovation is multi-core CPUs, (Central Processing Units) such as Intel® Core Duo® and Dual-Core Xeon® processors. Multi-core CPUs pack the processing power of several processors into a single chip. We are now at the point where standards-based general-purpose CPUs can provide the required processing power for many high-end imaging applications.

Higher performance processors are only part of the solution for high performance image processing. In order to transfer large amounts of data to and from the processors, high speed interfaces are also required. On this front, high-speed, bi-directional, serial differential I/O interfaces such as PCI Express (PCIe), Serial ATA (SATA), Serial Attached SCSI (SAS), and DDR2 plus fully buffered DIMM memory I/O interfaces provide great performance increases over their predecessors. These new standard high-speed interfaces can successfully compete with proprietary high-speed interfaces.

These multi-front technological advancements are drastically improving the ability of standards-based compute platforms to address the computational and I/O bandwidth challenges of high-end imaging applications.

A key advantage of standards-based systems is high manufacturing volumes that drive much lower costs than for proprietary systems. Also they are supported by a large and growing ecosystem of hardware and software manufacturers resulting in a tremendous pool of hardware and software designers competing with each other to constantly improve performance. By following established standards, these platforms inherently guarantee application software and firmware portability across platform generations and across semiconductor suppliers.

To understand the enormous advances in I/O bandwidth, Table 2 compares PCIe with PCI-X data transfer speeds.

COMPUTE COMPLEXITY OF IMAGING APPLICATIONS

High-end digital imaging often has a 2-dimensional FFT at its core. The computational overhead N in terms of MAC (Multiply-and-Accumulate) operations, of an n-point FFT is given as:

$$N = n\log2(n)$$

Hence, a 1024x1024 FFT of a 2D image has a computational overhead of 20*220, or about 21 million MACs. Additional operations and multi-channel scanning can easily cause this number to quickly rise to over 1 billion MACs. Floating-point additions are normally performed in 4 steps (compare exponents, equalize exponents, add mantissas, normalize the result). Floating-point multiplications are typically performed in 2 steps (add exponents simultaneously with multiplying mantissas, normalize the result). Hence, a floating-point MAC operation typically requires 6 steps. This means that a 2GHz CPU typically performs a single floating-point operation in 3x10-9 seconds. Hence, 1 billion MACs are completed in 3 seconds. When adding system memory and hard disk read/write overhead, the processing time for a single scanned slice can easily run into 30 seconds or more.

Even today's fastest mainstream general purpose dual-core CPUs require this time to be doubled when considering that operating systems and other applications, such as communication stacks snoop away cycles. High resolution and 3D images are produced by combining many 2D images or "slices" to build up the final image. This means that a multi-slice CT scan can take in excess of an hour to process on a single general-purpose host processor.

PCIe LINK	FULL-DUPLEX DATA RATE	PCI-X	DATA RATE
X1	5Gbps	66MHz	4.2Gbps
X2	10Gbps	100MHz	6.4Gbps
X4	20Gbps	133MHz	8.5Gbps
X8	40Gbps	266MHz	17Gbps
X16	80Gbps	533MHz	34Gbps
X32	160Gbps	N/A	N/A

Table 2. Data transfer speeds of PCIe compared to PCI-X (Actual data transfer rates will be lower due to protocol overhead)

MULTI-PROCESSING APPROACHES

Although a single general purpose processor can't provide the required processing power for high-speed image processing, neither can a single proprietary processor. The key is to split up the problem and perform the processing using a parallel approach.

One approach is to design embedded server cluster configurations, where the CPUs are connected to each other through a high-speed switching fabric, such as gigabit Ethernet or switched PCI-Express (PCIe). This approach creates a homogenous multi-processor configuration, where any CPU can become the controlling host computer. Multi-core processors fit into this category as well as they are essentially multiple processors combined on a single chip. A simple example of this approach is illustrated in Figure 1.

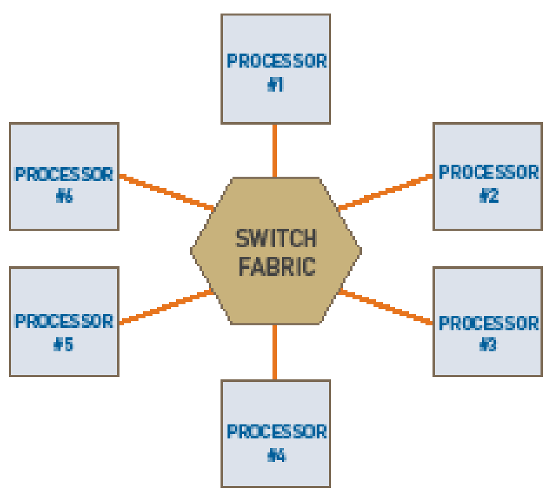

Figure 1. Homogenous multi-processor architecture

As with many things in life, there is no one-size-fits-all, "silver bullet" processor that is best for all imaging applications. While standards-based general purpose processors offer ever-increasing performance and reasonable costs, a system based on general purpose processors alone doesn't necessarily provide the optimal solution for all image processing applications. As a result, it is quite common to see a combination of general purpose processors and application-specific hardware, (hardware accelerators), particularly in high-performance systems. Hardware accelerators often provide very high processing performance. Unfortunately, they usually do so at the expense of programming flexibility. By using a combination of processor types, the system designer attempts to achieve a balance between processing performance, cost and programming flexibility. This type of system is often referred as "heterogeneous multi-processing" since the system contains multiple processors of more than one type.

One architectural approach for heterogeneous systems is to attach a number of application-specific hardware accelerators to an embedded server host computer. In this case the accelerators can be either attached directly to the host CPU's PCIe root complex, or via a high speed switch fabric. Figure 2 illustrates the heterogeneous multi-processing architecture approach.

A combination of the previous two designs results in a hierarchical, heterogeneous, multi-processing architecture. See Figure 4 for an illustration.

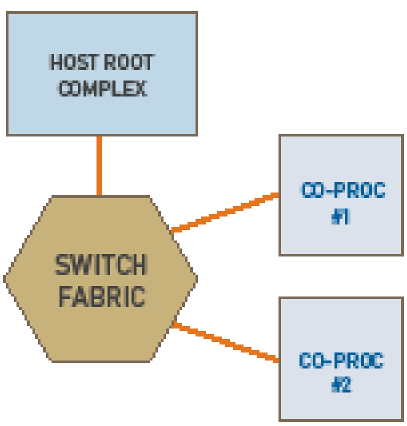

Figure 2. Heterogeneous multi-processor architecture

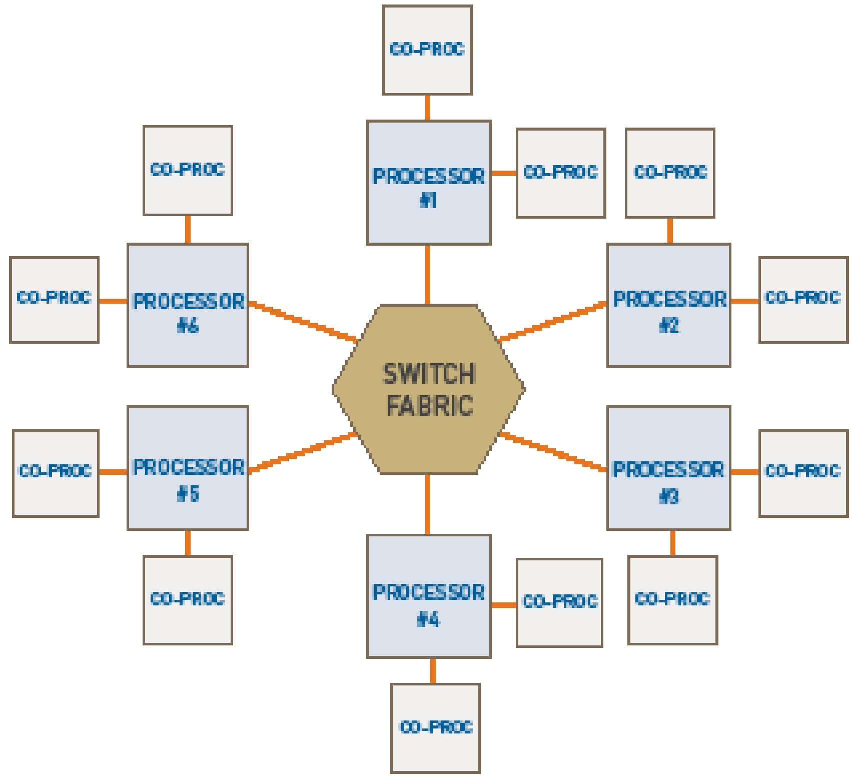

Figure 3. Hierarchical, heterogeneous, multi-processing architecture

HARDWARE ACCELERATORS (CO-PROCESSOR OPTIONS)

As mentioned above, high-end image processing was traditionally done in proprietary hardware. This often included ASICs, (Application-Specific Integrated Circuits). ASICs are custom chips that can be highly optimized for a specific application, but require large volumes to be cost-effective and also require long development times and significant up-front investment to produce the chips.

While ASICs are used in many image processing systems today, they are being replaced by programmable general-purpose hardware accelerators in most new designs. A popular choice of accelerator in imaging systems is an FPGA, (Field-Programmable Gate Array). FPGAs come in a variety of performance/price grades and often give the best "bang-for-the-buck" in terms of the image processing power than can be achieved for a given component cost. They are much more flexible than ASICs as they can be reprogrammed to meet changing needs, incorporate improved algorithms, etc. FPGAs are also cost effective at low to moderate volumes as the chips themselves are standard parts that are sold to a variety of customers across many industries. Another selling point is that FPGAs are relatively low power devices with typical TDP (Total Dissipated Power) of no more than about 40 Watts, even for very large FPGA chips. The downside to FPGAs is that programming of these devices is rather time consuming, and the programming model is less flexible than software for general purpose processors. The additional programming effort adds development cost and limits time-to-market.

A second type of accelerator that is gaining ground is the GPU (Graphics Processing Unit) - for instance NVidia® or ATI® graphics cards, which are programmable pixel processing engines. GPUs generally have strong floating point processing capability and very closely coupled high-speed memory, both of which are well suited for image processing. There are programming standards for GPUs that can provide similar benefits in terms of software portability as for general purpose CPUs. Additionally, many imaging systems already have a GPU card installed to support multiple and/or high resolution displays. In these cases, using spare GPU cycles to offload some image processing tasks from a general purpose CPU results in improved performance with no additional hardware cost. GPUs do come with some disadvantages, such as a different programming model than the host CPU and significant power dissipation, often 100 Watts or more.

There are quite a number of additional hardware accelerators on the market that handle tasks such as parallel processing very efficiently and could offer significant benefits for image processing applications. Time will tell which ones achieve market success.

There are a number of ways that heterogeneous systems can be architected. One way that warrants mention is abstracting the hardware accelerators in software via API calls so that they function as co-processors for a general purpose processor. In this case, all of the application code can be run in software on the general purpose processor for the initial implementation. Specific functions are then chosen to be handed-off to the co-processor(s) for acceleration. A key advantage to this approach is that the application software that runs on the x86 host processor remains portable. Additionally, this type of system can be scaled to achieve various price/performance targets. A low end 'base" system might use only the general purpose processor. This same system could be configured for higher performance, and command a higher price, by simply adding one or more co-processors that accelerate specific compute-intensive functions.

Another potential advantage of the co-processor approach is that the image processing tasks can be partitioned in ways that allow for migration of larger portions to the general-purpose host processor as increases in processing capacity are achieved with each new generation. This means that imaging system manufacturers can create a road map for their products in which today's high-end and expensive systems are replaced by tomorrow's inexpensive, general-purpose systems and new higher performance systems can be developed.

BUILDING ON INDUSTRY STANDARDS

Using mainstream processors and operating systems is one part of the cost-down equation of the homogeneous and heterogeneous architectures of the previous section. The other part of that equation is the use of mainstream system-level infrastructure.

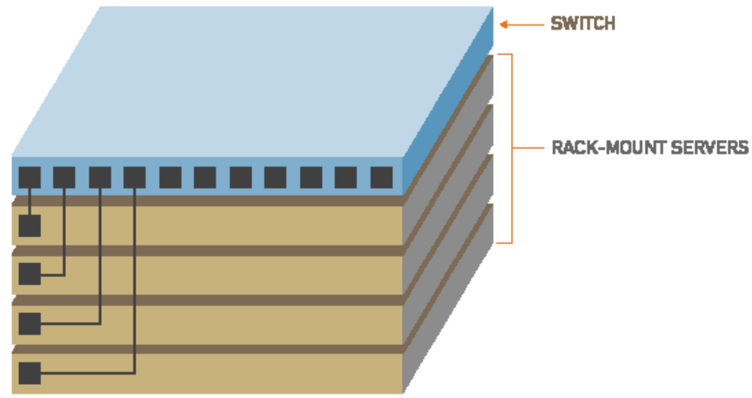

Figure 4. Hierarchical, heterogeneous, multi-processing architecture

There are at least a dozen or so mainstream system-level standards to choose from for an embedded server. However, the most relevant ones to high-end digital image processing systems are:

- Standard motherboards in the ATX or SSI form factors
- The COM Express Basic and Extended form factors.

Each of these form factors further benefits from standardized mechanical enclosures. This in turn assures that high-end digital image processing systems can be built cost-effectively with a large pool of standard components and suppliers to choose from.

Note: The COM Express form factors do not have a specific enclosure defined, but are system-level sub-components that can easily housed in the enclosures defined for the other standard form factors.

ATX AND SSI

A high-end digital image processing system built with ATX or SSI motherboards in standardized 19" rack-mount 1U, 2U, 3U or 4U* enclosures can easily be scaled from single systems to clusters. Each node in the cluster can accommodate one or more co-processors in the shape of add-on cards that communicate with the host's PCIe root complex. Each node in the cluster can be connected to an Ethernet switch fabric, with the switch hosted in one of the nodes or externally. Such configurations create a hierarchical multi-processing system of localized multi-processing at each node and global multi-processing through the switch fabric. See Figure 4. For more information on the ATX form factor see www.formfactors.org. For information about the SSI form factor, see www.intel.com.

** Note: "U" is an abbreviation for a "rack-unit" of height. (1U = 1.75 inches)*

COM Express. The COM Express (COM-E) specification defines a modular host system or computer-on-module (COM), where each module can be a fully autonomous host system. A COM-E module provides:

- Support for a single or multi-core CPU
- High-speed memory
- PCIe links up to x32 wide
- SATA, SAS and 1Gigabit Ethernet interfaces
- Provisions for 10Gigabit Ethernet and future speeds of PCIe, SATA and SAS

COM Express modules allow OEMs to scale their systems from floor-installed to bench-top, to portable and handheld systems, using COM Express as the sole platform architecture. COM Express does not specify a particular chassis, but due to the component-like nature of the modules, any of the afore-mentioned chassis can be used. COM Express modules attach to a motherboard or blade via mezzanine connectors. See Figure 5. For more information about the COM Express specification see www.picmg.com.

COM Express provides a practical, standards-based platform for the imaging applications described here.

SOLUTIONS PROVIDED BY RADISYS

RadiSys provides innovative compute solutions for imaging applications that address all the different approaches presented above. Our solutions are built around industry standards, featuring Intel® x86 processors including the latest multi-core designs.

Standardization and the use of mainstream parts can dramatically lower cost, as one benefits from the economies of scale for such components. Yet, it is important to carefully select among these, using a set of criteria that optimizes price, performance and longevity of the solution. RadiSys specializes in long-life solutions with typical production lifetimes of 5 years. We utilize components specifically selected for extended duty cycles in industrial environments. Our solutions are designed with low power dissipation, low acoustic noise and low electromagnetic emission requirements in mind.

Our Procelerant™ RMS420-5000XI embedded server is an example of a high performance product developed for imaging applications. The RMS420-5000XI features two Intel® Dual-Core Xeon processors (Quad-core coming in 2007). These processors incorporate the new Core® architecture from Intel that provides superior performance for SIMD (Single-Instruction, Multiple Data) parallel processing that is at the heart of most image processing applications. The RMS420-5000XI also offers PCIe x16, x8 and x4 connectivity to support GPU cards, co-processors and/or high speed data I/O. Dual on-board gigabit Ethernet ports provide additional high-speed I/O capability.

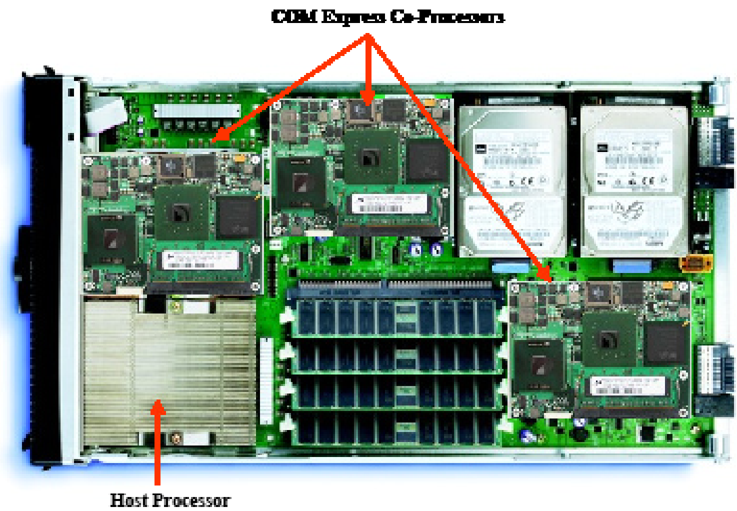

Figure 5. COM-E Mezzanine Modules on a Host Platform

Our Procelerant™ CE line of COM Express modules includes both mono- and dual-core processor products, based on the high-performance and power-efficient Intel® Core Solo™ and Core Duo™ processors. Uniquely, we offer up to 4GB of dual-channel memory on "Basic" size COM-E modules. This feature is critical for high performance image processing in a small form factor.

CONCLUSIONS

High-end imaging systems are moving from proprietary hardware to standards-based solutions. System manufacturers are seeing benefits in terms of increasing performance, reduced costs for both systems and development, and faster time-to-market. Heterogeneous systems that combine a general purpose processor with one or more specialized co-processors often offer the best solution for challenging high-end image processing applications.

RadiSys Corporation manufactures standards-based systems for image processing that provide top performance coupled with product long-life, and meet imaging customer needs for low EMI, low acoustic noise, small form factors and low power dissipation.

Our Procelerant embedded server and COM Express product lines provide the performance, scalability, and flexibility required for high-end digital imaging systems. To learn more about our products visit www.radisys.com.

RadiSys
THE POWER OF WE

World Headquarters
5445 NE Dawson Creek Dr
Hillsboro, OR 97124 USA
Phone: 503-615-1100
Fax: 503-615-1121
Toll-Free: 800-950-0044
www.radisys.com
info@radisys.com