Standard Guide for
Hardware Implementation for Computerized Systems[1]

1. Scope

1.1 This guide provides assistance in the choice of computing hardware resources for ship and marine environments and describes:

1.1.1 The core characteristics of interoperable systems that can be incorporated into accepted concepts such as the Open System Interconnection (OSI) model;

1.1.2 Process-based models, such as the Technical Reference Model (TRM), that rely on interoperable computing hardware resources to provide the connection between the operator, network, application, and information; and,

1.1.3 The integrated architecture that can be used to meet minimum information processing requirements for ship and marine environments.

1.2 The use of models such as OSI and TRM provide a structured method for design and implementation of practical shipboard information processing systems and provides planners and architects with a roadmap that can be easily understood and conveyed to implementers. The use of such models permit functional capabilities to be embodied within concrete systems and equipment.

1.3 The information provided in this guide is understood to represent a set of concepts and technologies that have, over time, evolved into accepted standards that are proven in various functional applications. However, the one universal notion that still remains from the earliest days of information processing is that technological change is inevitable. Accordingly, the user of this guide must understand that such progress may rapidly invalidate or supersede the information contained herein. Nonetheless, the concept of implementing ship and marine computing systems based on these functional principles allows for logical and rational development and provides a sound process for eventual upgrade and improvement.

2. Referenced Documents

2.1 *ASTM Standards:*[2]
E1013 Terminology Relating to Computerized Systems (Withdrawn 2000)[3]
F1757 Guide for Digital Communication Protocols for Computerized Systems

2.2 *ANSI Standards:*[4]
X3.131 Information Systems—Small Computer Systems Interface-2 (SCSI-2)
X3.172 American National Standard Dictionary for Information Systems
X3.230 Information Systems—Fibre Channel—Physical and Signaling Interface (FC-PH)
X3.232 Information Technology—SCSI-2 Common Access Method Transport and SCSI Interface Module
X3.253 Information Systems—SCSI-3 Parallel Interface (SPI)
X3.269 Information Technology—Fibre Channel Protocol for SCSI
X3.270 Information Technology—SCSI-3 Architecture Model (SAM)
X3.276 Information Technology—SCSI-3 Controller Commands (SCC)
X3.277 Information Technology—SCSI-3 Fast-20
X3.292 Information Technology—SCSI-3 Interlocked Protocol (SIP)
X3.294 Information Technology—Serial Storage Architecture—SCSI-2 Protocol (SSA-S2P)
X3.297 Information Systems—Fibre Channel—Physical and Signaling Interface-2 (FC-PH2)
X3.301 Information Technology—SCSI-3 Primary Commands (SPC)
X3.304 Information Technology—SCSI-3 Multimedia Commands (MMC)
MS58 Information Technology—Standard Recommended Practice for Implementation of Small Computer Systems

---

Interface (SCSI-2), (X3.131.1994) for Scanners

NCITS 306 Information Technology—Serial Storage Architecture—SCSI-3 Protocol (SSA-S3P)

NCITS 309 Information Technology—SCSI-3 Block Commands (SBC)

2.3 *IEEE Standards:*[5]

100 Standard Dictionary for Electrical and Electronic Terms

488 Digital Interface for Programmable Instrumentation

610.7 Standard Glossary for Computer Networking Terminology

796 Microcomputer System Bus

802.11 Wireless LAN Medium Access Control and Physical Layer Specifications

1003.2d POSIX—Part 2 Shell and Utilities—Amendment: Batch Environment

1003.5 Binding for System Application Program Interface (API)

1003.b Binding for System Application Programming Interface (API)—Amendment 1: Real-time Extensions

1014 Versatile Backplane Bus: VMEbus

1101.10 Additional Mechanical Specifications for Microcomputers using the IEEE Std 1101.1 Equipment Practice

1155 VMEbus Extensions for Instrumentation: VXIbus

1212.1 Communicating Among Processors and Peripherals Using Shared Memory (Direct Memory Access DMA)

1394 High Performance Serial Bus

1496 Chip and Module Interconnect Bus: Sbus

1394 32-bit Microprocessor Architecture

2.4 *ISO Standards:*[4]

1155 Portable Operating System Interface for Computer Environments (POSIX)

9945-1 System Application Program Interface (API) [C language]

9945-2 Shell and Utilities

2.5 *TIA/EIA Standard:*[6]

568-A Commercial Building Telecommunications Cabling Standard

3. Significance and Use

3.1 This guide is aimed at providing a general understanding of the various types of hardware devices that form the core of information processing systems for ship and marine use. Ship and marine information processing systems require specific devices in order to perform automated tasks in a specialized environment. In addition to providing information services for each individual installation, these devices are often networked and are capable of supplementary functions that benefits ship and marine operations.

3.2 A variety of choices exists for deployment of information processing devices and greatly increases the complexity of the selection task for ship and marine systems. The choice of a particular device or system cannot be made solely on the singular requirements of one application or function. Modern

information processing systems are usually installed in a complex environment where systems must be made to interact with each other. Ship and marine installations add an even further layer of complexity to the process of choosing adequate computerized systems. This guide aims to alleviate this task by giving users specific choices that are proven technologies that perform in a complex environment.

3.3 Hardware resources used in ship and marine installations are a result of careful consideration of utility and function. These resources may require some physical specialization in order to inhabit a particular environment, but they are in no way different from equipment used in shore-based situations. Ship and marine computer system configurations, interconnections, and support services are essentially the same as those found in a land-based network environment and as a result, the skill sets of ship and marine information processing system users, administrators, and support personnel are interchangeable with those of shore-based activities.

4. Standards Profiles

4.1 Standards profiles are sets of specifications bundled together to describe the technical standard for a function or a service (such as operating systems, network, and data interchange services), and will include minimum criteria for the information and technology that support specific functional requirements. Profiles equate to the lowest level process, and document agreed-to implementation requirements used in building and operating systems. Systems using the same standards, but different options, will probably not interface correctly. The Technical Reference Model (TRM) is useful for assembling standards profiles across technology categories of Computing Resources, Information Management, and Applications.

4.1.1 The TRM identifies and specifies the support services (multimedia, communications, and so forth) and interfaces that provide a common operating environment and support the flow of information among enterprise and common support applications. This model represents the computer resources, information management, and applications categories and interfaces with the communication and networking technology categories that are appropriately represented by the ISO Open System Interconnect model. The TRM addresses standard profiles that provide seamless application support over widely distributed computing resources and attendant interfaces between the computing resources and other technologies.

4.2 Computing hardware resources represent generally consists of Central Processing Unit(s) (CPU), Input and Output (I/O) interfaces, main memory, buses, and peripherals. The external environment considerations that affect computing hardware resource selection are security, communications, real-time, and high availability. The computing hardware resource provides the environment necessary to support application software. From the perspective of the application software, services are provided by the computing resource, whether the particular services are provided locally or remotely as part of a distributed system.

4.3 The architecture needed to support a typical application consists of computers that perform as clients and servers. The

servers host the primary application software and contain the processing power necessary to support multiple users. Servers also host the data needed to support the application. The standard 3-tiered application architecture consists of (*1*) an application server, (*2*) a data server, and (*3*) presentation clients (see Fig. 1).

4.4 In the future, most application processing software will be hosted on the server computers. Clients will use presentation software that connects to the servers using a common interface. At that time, client computers will likely be less expensive and tailored to the user's individual preference because application interoperability will not be a significant factor.

4.5 Today, however, most application software is hosted on the client and interoperability among clients is a critical factor. Even within the client-server application architecture, application specific software resident on the client is still prevalent. This demands consistency of client workstations across an entire installation to achieve seamless interoperability. Table 1 outlines a rationale for the client-server deployment strategy.

4.6 Driven by the current state of client-server technology, the general philosophy for implementing computing resources is the concept of homogeneous clients and heterogeneous servers. Homogeneous clients facilitate providing a consistent interface between the user and the system and make system support and maintenance less complex. Heterogeneous servers support the various computing requirements of applications needed to support ship and marine operations. The same advantages that homogeneous clients enjoy can be achieved if servers are homogeneous as well. Independent of whether or not the server suite employed is heterogeneous or homogeneous, it is important that they perform their function transparently to the user (that is, the user neither knows nor cares about the location, number, or vendor of the server being used.) Requiring servers to be homogeneous would restrict the introduction of new server technology, choking innovation and preventing the installation from taking advantage of advances in computing such as massively parallel processors.

5. Computing Hardware

5.1 *Computing Resources*—Computing resources consist of many computing hardware components and configurations of these components. This section covers the various hardware components that make up a computing resource system and examines how these components are commonly configured.

5.2 *Component Technologies*—The major hardware components of Computing Resources are the Central Processing Unit (CPU), one or more backplane buses, main memory (both RAM and cache), Input/Output (I/O) interfaces, and peripherals. This section will examine each of these areas and provide guidance on the selection of these component technologies as part a computing resource system.

5.2.1 *CPU*—The CPU is the "engine" of the computer system and, combined with the OS (operating system), forms the core of the computing resource. Since the OS drives many decisions concerning the computer resource, a CPU that is compatible with the OS becomes an overriding factor in determining the type of CPU. Other than the OS, the main factors to consider in determining the type of CPU for the computer are processing speed (performance) and cost. For computing resources, such as servers and multiprocessors, scalability of the number of processors can be a significant factor in determining CPU.

5.2.2 *Bus*—The computer bus connects the different components of the computer resource together and allows them to pass data between them at high speeds. Computer resource configurations, such as personal workstations, often limit or determine the type of bus that will be used. Often there are


FIG. 1 Three-Tiered Application Architecture

TABLE 1 Client-Server Deployment Rationale

Rationale for Heterogeneous Servers	The server must be tailored to the specific application that may not be supportable by computers most prevalent in the marketplace.
	Many applications work well in their current computing environment and it is not cost effective to change.
	It is not practical to have all applications on a common server for multiple reasons including the need to maintain competition between computer developers and vendors.
	Encourages innovation by not restricting the type of computer used for the development of applications.
Rationale for Homogeneous Clients	Allows for a common, consistent user interface.
	Maximizes interoperability.
	Minimizes re-training required as users transfer to different organizations within the enterprise.
	Maximizes the ability to use common support and maintenance skills, parts and labor; thereby minimizing cost.
	Maximizes portability of support for applications across the enterprise as well as portability of user skills.
	Allows for economies of scale in both procurement (volume discounts) and support (more focused skill set for help desk personnel).

TABLE 2 CPU Sample Implementations

Clients, Servers, and Special Purpose
Intel Pentium/Celeron
AMD K6/Athlon/Duron
VIA C3/Cyrix
Motorola PowerPC
Transmeta Crusoe

Servers and Special Purpose
MIPS 32/64
Compaq Alpha
Hewlett Packard PA-RISC
Sun Microsystems SPARC/UltraSPARC
Motorola PowerPC

multiple buses connected together to allow for multiple types of component cards or to extend a non-expandable system bus. Considerations in determining the type of bus to use are: number and type of commercial products compatible with the bus architecture, number of parallel data bit lines, clock speed, and cost. Once the appropriate bus architecture is determined, an important computer resource factor becomes how many interface slots are available on the bus for component cards.

TABLE 3 Sample Bus Implementations

PCI
VME/VXI
SBUS
VXI
CardBus
GPIB/HP-IB
EISA

Note 1—Peripheral Connect Interface (PCI) bus is quickly gaining favor as a low-cost preferred system bus architecture. PCI provides the necessary throughput to support the high-end data rates required by many of today's applications. Most commercially available computers come with a PCI bus.

Note 2—EISA bus should be used only to accommodate legacy systems.

Note 3—Although VME bus is more popular, VXI bus offers a greater degree of standardization and therefore a greater degree of interoperability between vendor's products. These are the buses of choice for embedded systems. IEEE Std 1014 applies to VME and Std 1155 applies to VXI.

Note 4—GBIB is the Standardized version (IEEE Std 488) of the HP-IB implementation developed by Hewlett-Packard.

Note 5—CardBus is the new 32-bit high performance bus defined by the new PC Card Standard released by the PCMCIA standards body and trade association. The PC Card Standard replaces the outdated PCMCIA version 2.0 and version 2.1 standards.

5.2.2.1 Use buses that provide the necessary performance economically and are compatible with the board level components that are needed to meet requirements. For buses that provide slots for component cards, use standard buses that are supported by multiple vendors providing compatible component cards.

5.2.3 *Main Memory*—Main memory is the storage warehouse of the computer where data and programs are stored for efficient processing. In the context of this section, main memory refers to cache and RAM. The main factor to consider in acquisition of a computer system is the quantity (in megabytes) of RAM. Other considerations are access speed, mounting design, and parity. Computer systems with too little memory run slowly, won't load, and crash often. Mounting designs today generally provide for easily upgradeable Memory Modules. SDRAM (Synchronous Dynamic Random Access Memory) has long been a standard for memory, but more advanced designs such as DDR (Double Data Rate) and Rambus memory offer better speed and throughput and are rapidly gaining wide acceptance. Older architecture memory designs are generally slower and less efficient and should be avoided to the extent possible.

5.2.3.1 Cache is usually hard wired to the motherboard and has a faster access time than RAM. Computer system caches of 512 KB or larger are generally satisfactory.

5.2.3.2 RAM can often be on a separate memory board and is used to store the OS, applications that are running, and data files. The amount of RAM needed for a computer system can vary with the environment and the OS. Servers generally need about an order of magnitude more RAM than personal workstations. For personal workstations with a 16/32 bit operating system, 64 MB or more of RAM is recommended; for workstations with 32/64 bit operating systems, 128 MB or more is recommended. For servers, use the Network Operating System (NOS) guidelines based on the environment. The three major factors used to determine the amount of RAM for a server are number of user connections, number of processes running, and amount of hard drive space.

5.2.4 *Input/Output (I/O) Interfaces*—I/O interfaces allow the computing resource to move data between the "outside world" and the CPU and main memory. Operations like loading a program or file from a floppy or CD, sending and receiving information over the LAN or WAN, or sending a document to the printer to get a hardcopy use I/O interfaces.

TABLE 4 Recommended Memory (RAM) Standards

Dynamic Random Access Memory (DRAM)
Extended Data Output Dynamic Random Access Memory (EDO DRAM)
Synchronous Dynamic Random Access Memory (SDRAM)
Double Data Rate Memory (DDR)
RAMBUS Memory

NOTE 1—Extended Data Out (EDO) DRAM was long a standard for mass-market memory, but Synchronous DRAM (SDRAM) is now the standard for currently installed machines. SDRAM is a memory architecture that incorporates many improvements over traditional DRAM technologies. SDRAM is a new technology that runs at the same clock speed as the microprocessor. The clock on the memory chip is coordinated with the clock of the CPU, so the timing of both the memory and the CPU are "in synch." This reduces or eliminates wait states, and makes SDRAM significantly more efficient than Fast Page Mode (FPM) or even Extended Data Out (EDO) memory.

NOTE 2—Even more speed and throughput improvements are being realized with Double Data Rate (DDR) and Rambus Memory; selection of a memory architecture will need to be made according to a careful consideration of cost (particularly over the anticipated service life of a system) and performance considerations.

Quite often the information is sent to or received from a peripheral, which is discussed in the next section.

5.2.4.1 Use I/O interfaces that use open access standards, support open device connections, and are platform independent.

5.2.5 *Peripherals*—Peripherals provide data access, input, storage, and connectivity for a computing resource. The number of peripherals available on the commercial market continues to explode, generally driven by processor speeds, memory/storage capacities, and I/O speeds. Although there are many different types of peripherals, such as printers, facsimiles, modems, scanners, video cameras, microphones, speakers, and so forth, the main issue in specifying/procuring these items is the compatibility of their I/O interfaces with the computer (see 5.2.4) and application software. Apart from these compatibility issues, the major considerations for acquiring peripherals are cost and performance (which can include both speed and quality). A major category of peripherals is static storage devices. As distinguished from main memory (covered in 5.2.3), static storage devices retain data when the power is off. The remainder of this section will discuss storage devices. Use peripherals that support standard I/O interfaces and are platform independent.

5.2.5.1 *Storage Device Standards*—Storage refers to the capability to store information outside the central processor.

TABLE 5 Recommended I/O Interface Standards

SCSI-2	SCSI-3	FC-PH/FC-AL	ESCON	IPI	PC Card
USB	IEEE 1394	XIO	Serial	Parallel	

NOTE 1—Fibre Channel (FC-PH) is emerging as a host-level interface standard for delivery of high I/O data transfers. FC-PH provides connections for workstation clustering, storage clustering and network-based storage concepts, parallel processing, load leveling, host-to-host or server-to-server communications, host- or server-to-mass storage communications, bulk data transfer, and multimedia. FC-PH is also being used as a system bus at the CPU and memory level as well as a way to cluster multiple systems similar to Non-Uniform Memory Access (NUMA). NUMA describes an architectural approach to clustering multiple systems such that distributed memory appears to the operating system as shared memory. This architectural approach allows a benefit to the user of a shared memory programming model with the scalability of a massively parallel processor's (MPPs) distributed memory model. FC-PH currently provides up to 1.062 GigaBits Per Second (Gbps) bandwidth using optical fiber.

NOTE 2—FC-PH-based storage subsystems will be supported on mainframes; therefore, the Enterprise Systems Connection or ESCON-based systems will be phased out.

NOTE 3—The Upper Layer Protocol (ULP) over FC-PH from the mainframe will initially deploy as SCSI-2 or SCSI-3 and migrate quickly to Intelligent Peripheral Interface (IPI). IPI will be deployed on large servers before it will be deployed on mainframes. Note that IPI is a protocol only; SCSI is both a protocol and an interface. FC-PH is an interface that can support various ULPs.

NOTE 4—For servers and higher end personal workstations, SCSI is the predominant host-level interface for current disk and peripheral devices. As the demand for higher data transfer rates and expanded connectivity requirements increase, SCSI-2 and SCSI-3 over FC-PH in a PCI form factor will be the appropriate long-term direction.

NOTE 5—Enhanced IDE is the predominant peripheral interface for lower end personal workstations. SCSI is the current interface of choice for high-speed devices, such as high-capacity disk and tape drives and is the appropriate long-term direction for both the commodity level and high capacity devices.

NOTE 6—Personal Computer Memory Card International Association (PCMCIA) announced that PCMCIA cards are now referred to as PC Cards. The PC Card Standard defines a 68-pin interface between the peripheral card and the PC Card "socket" into which it gets inserted. It also defines three standard PC Card sizes, Type I, Type II and Type III. All PC Cards measure the same length and width, roughly the size of a credit card. Where they differ is in their thickness. Type I, the smallest form factor, often used for memory cards, measures 3.3 mm in thickness. Type II, available for those peripherals requiring taller components such as LAN cards and modems, measures 5 mm thick. Type III is the tallest form factor and measures 10.5 mm thick. Type III PC cards can support small rotating disks and other tall components.

NOTE 7—As mentioned under "Bus" (see 1.1.2), CardBus is a new standard introduced as an addendum to the PCMCIA PC Card standard. CardBus is a 32-bit bus-mastering card operating at 33 Mhz transferring data at up to 132 MBytes per second. (The 16-bit PC Card bus data rate is 20 MBytes per second.) Like PC Cards, CardBus uses the 68-pin interface but operates at 3.3 volts versus the 5 volts used by PC Cards. CardBus slots are backward compatible with PC Cards. Card sockets can support PC Cards only or CardBus cards only, but not a mixture of the two.

NOTE 8—Serial I/O refers to standard serial interfaces such as RS 232, 422, 423, and 449.

NOTE 9—Parallel I/O refers to standard parallel interfaces such as micro-Centronix.

NOTE 10—XIO is an emerging I/O solution based on Specialix' SI controller, and incorporating the high performance I/O processor and communication technology introduced in the RIO range. XIO is ideal for the 16-32 user system where sustaining high performance is essential.

NOTE 11—IEEE 1394, also known as "Firewire," is a high speed serial I/O that supports data rates up to 400Mbps.

For most computers, the predominant technology for storage has been magnetic disk and will remain so for the next few years.

(1) Storage media is the physical material on which data is stored. The choice of media is usually determined by the application needs in terms of data accessibility, storage density, transfer rates, and reliability. Broad industry standards exist for 3.5-in. magnetic disk, 4-mm digital audio tape (DAT), 8-mm helical tape, digital linear tape (DLT), ½-in. tape, and compact disk-read only memory (CD-ROM). However, industry is in the process of agreeing on a standard for Digital Versatile Disk (DVD) which will able to read CD-ROMs as well as the new DVDs. Standards do exist for write once read many (WORM) optical and magneto-optical (MO) disks. Although the physical medium for optical technologies conforms to an open standard, the device's recording format may not.

(2) Archived data on outdated storage media (for example, 5.25-in. floppy disks) should be transferred to media that is more current to avoid data being "trapped" on obsolete media that cannot be read by devices currently on the market.

(3) As networks proliferate and storage requirements expand, storage technology that uses standard interfaces and promotes hardware and software supplier independence is necessary. This technology will enable us to take advantage of the open systems environment.

(4) Implement storage technology and storage device media that use open access standards, support open device interface standards, and are platform independent.

5.2.5.2 *RAID Technology*—Redundant Array of Independent Disks (RAID) technology protects from data loss by providing a level of redundancy immediately within an array. The array contains removable disk drive modules that are automatically rebuilt in the event of a device failure without causing the system to shut down. When RAID levels other than 0 are used, no downtime is required to replace a failed disk drive. Data is continuously available while reconstruction of the failed disk occurs in the background. Much of the benefit of RAID technology lies in its capability to off-load storage management overhead from the host system. To realize this benefit, RAID developers endow their array controllers with significant levels of intelligence. For instance, Adaptive RAID supports multiple RAID levels based on workload characteristics. Choose the RAID level based on your specific need.

5.3 *System Configurations*—The hardware component technologies mentioned in 5.2 can be configured in many different ways to accomplish different tasks and meet different requirements. This section examines some of the common configurations (Personal Workstations, Servers and Embedded Computers) with guidance on what component technologies to use for each configuration.

5.3.1 *Personal Workstation*—Personal Workstations (PW) are devices that contain at least one CPU (sometimes several) and provide a user interface, typically a GUI, as well as personal productivity tools, local data storage, and a flexible method for accessing and manipulating data. These PWs are commonly known as Personal Computers or PC's, desktop computers, portables (laptops or notebooks), or workstations and use one of several bus architectures. Low-end PWs are

TABLE 6 Storage Device Media

3.5 in. magnetic disk
3.5 in. floppy disk
5.25 in. WORM
5.25 in. MO
ISO 13346
CD-ROM
Flash Memory Removable Devices
Smart cards
DVD
½-in. helical tape
4 mm DAT
DLT Tape
3490E cartridge tape

Note 1—8-mm helical tape is appropriate for backup and archive use; however, it is slow and unreliable for near-line storage solutions where frequent access is required. 4mm DAT offers a faster, more reliable solution to high-capacity and high-access applications.

Note 2—½-in. helical tape offers higher storage capacity than 4 mm or 8 mm; however, it is expensive, proprietary, and only Storage Technology supports it.

Note 3—3490E square cartridge ½-in. tape is predominantly a mainframe technology for off-line data storage.

Note 4—High capacity tape applies capacity multipliers of 2× to 4× to current technology. Gains are achieved through increased density factors as well as media length (number of media ft per unit).

Note 5—Flash Memory Storage devices are finding wide acceptance in Asia. They can hold as much as 128 MB of data in a keychain form factor that plugs into an available USB port.

Note 6—Smart cards are being used in Europe and Asia. Smart cards will be used predominantly as an intelligent storage media for providing services to individuals.

Note 7—The Digital Video Disk (DVD, also known as Digital Versatile Disk) allows for dual-sided as well as dual-layered implementations that will increase CD-ROM capacities significantly. In addition, transfer rates may improve to as high as 16 Mega Bits Per Second (Mbps).

used primarily to support the general office work place. More powerful PWs are predominantly used in high-end computer applications such as Computer-Aided Design/Computer-Aided Manufacturing/Computer-Aided Engineering (CAD/CAM/CAE), application development, multimedia, and decision support data analysis presentation.

5.3.1.1 Also included in PWs are handheld computers—Personal Digital Assistants (PDAs), also referred to as Personal Information Managers (PIMs). Handhelds are computer systems that fit in a person's hand and are extremely portable. Handheld systems also tend to have two types of input methods: pen or keyboard. The pen input can be used as digital ink, a mouse, or for handwriting recognition.

5.3.1.2 Combine components that provide flexible, scaleable, and easy-to-use personal workstations that support the Client/Server model of computing, data access, and multimedia. Allow for an external communication device such as a modem or network interface card.

5.3.1.3 Table 9 provides a quick summary of the nominal specifications for Personal Workstation implementation.

5.3.2 *Servers*—Servers are computing resources that can be configured to support groups from small teams (work group servers) to entire ships (campus servers). Work group servers provide support, such as directory, file, print, and authentication services, in a LAN environment. Campus servers may augment and, for many new applications, replace traditional