Open standards for high availability and system management

by Deepak Shahane and Louise Moser, SA Forum

THE SERVICE AVAILABILITY FORUM IS A CONSORTIUM OF COMPUTER AND COMMUNICATIONS COMPANIES WORKING TOGETHER TO DEVELOP OPEN STANDARDS FOR HIGH AVAILABILITY AND SYSTEM MANAGEMENT SOFTWARE AND HARDWARE. THE SA FORUM DEVELOPS AND PUBLISHES INTERFACE SPECIFICATIONS, AND PROMOTES AND FACILITATES THE ADOPTION AND USE OF THE SPECIFICATIONS THAT IT DEVELOPS.

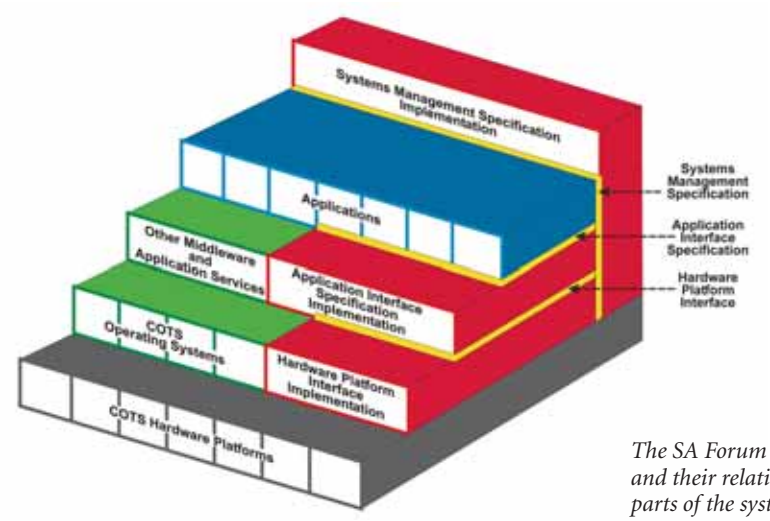

The SA Forum specifications and their relationships to other parts of the system

Modern integrated computer and communications systems provide multiple services: voice, data, management, security, VoIP, wireless, video conferencing, Web services, etc. Such systems are assembled from commercial-off-the-shelf (COTS) hardware and software components that are sourced from multiple vendors. Within the lifetime of such a system new services will be needed, hardware components will be replaced by new hardware components, possibly from different vendors, and software components will be enhanced by new software, probably programmed by different programmers.

Proprietary designs and interfaces make it difficult, if not impossible, to integrate components provided by different vendors and different programmers, into a single system. Therefore, modern integrated computer and communications systems need standardized interfaces and components, rather than the proprietary designs and infrastructure of the past. Today's computer and communications systems must provide continuous service for their users, even though hardware and software faults occur, and even while the hardware and software are being upgraded. The major telecommunications systems of the past were inordinately costly to develop and required years to reach deployment. Significant portions of those costs and timescales resulted from the need for high availability. Because high availability standards did not exist, high availability was provided by proprietary custom techniques within the application programs. Such custom coding, though successful, added considerable complexity to the application programs, increased the development cost and time, and required application programmers with special skills in high availability programming.

Industry standards for high availability reduce the cost and time of development, by ensuring that the application programs are not made more complex by custom high availability code that is intermingled with the application code. Industry standards allow the application programmers to focus on the application logic and the services it provides, and minimize the need for special high availability programming skills. They also allow new hardware and software components to be integrated together into computer and communications systems with greater ease.

The Service Availability Forum (SA Forum) is a consortium of computer and communications companies working together to develop open standards for high availability and system management software and hardware. The SA Forum develops and publishes interface specifications, and promotes and facilitates the adoption and use of the specifications that it develops. The SA Forum interface specifications enable the use of COTS software and hardware components to construct high availability systems and services that provide uninterrupted service for their users. The SA Forum has recently released new versions of the hardware platform interface (HPI) B.01.01, and the application interface specification (AIS) B.01.01. Currently, the SA Forum is working on a systems management specification. The Figure shows the SA Forum specifications and their relationships to other parts of the system.

The HPI is an interface between the application software and middleware, and the underlying hardware that allows portability of software across different hardware platforms. The HPI allows programmers of applications and middleware to write software that is independent of the particular hardware platform. The HPI specification represents the platform-specific characteristics of the physical hardware in an abstract model (with two different representations, a physical view and a management view), and defines standard function calls that support that model. A vendor's implementation of the HPI represents the physical hardware in terms of the abstract model, and translates the function calls defined by the specification into appropriate actions of the physical hardware.

The HPI allows the user to set and retrieve configuration or operational data about the hardware components, and to control the operation of those components, such as startup, shutdown and testing. It allows the user to subscribe to receive events and forwards events to the users that have subscribed for them. It provides hot-swap capabilities that allow a hardware component to be added to, or removed from, the system with little or no perturbation. An HPI implementation monitors and controls the physical hardware, and provides services to the applications and other middleware, independent of the hardware platform. It discovers the capabilities of the hardware platform, and maps those capabilities into the model, which the middleware maintains and presents to the applications and other middleware. It reads sensor values and operates watchdog timers to monitor the health of the hardware components.

The latest version of the HPI specification, HPI B.01.01, contains modifications to the original version, HPI A.01.01, that simplify the use of the HPI, add several new features and clarify items that adopters found ambiguous. For example, it improves the usability of sensor representations, updates the inventory data repository concept, clarifies the notion of HPI domains, adds support for reporting failed resources and introduces a new management instrument type.

The AIS defines an interface for high availability applications that is independent of different vendors' implementations of the high availability middleware. The AIS allows application programmers to write application software that is portable across different vendors' implementations of the high availability middleware. The AIS represents the high availability characteristics of the system in an abstract model, and defines APIs for functions that support that model. A vendor's implementation of the AIS represents the high availability middleware in terms of that abstract model, and translates the APIs defined by the specification into appropriate actions of the high availability middleware.

The AIS defines extensive APIs, which an application programmer can use in conjunction with a vendor's implementation of the AIS. It defines APIs for an availability management framework, including APIs for registration and de-registration, health monitoring, availability management, protection group management and error reporting. In addition, it defines APIs for a cluster membership service, checkpoint service, event service, message service and lock service. An AIS implementation monitors and controls the applications and the middleware, and detects and responds to faults. The availability management framework coordinates redundant resources within a cluster and enables highly available applications. It determines the readiness state and the HA State of a component, and checks the health of a component by invoking callback functions of the component. A component can, in turn, query the framework for information about other components' states.

The latest version of the AIS specification, AIS B.01.01, restructures the original version, AIS A.01.01, into seven volumes. Volume 1 presents the objectives of the AIS specification, the system model, the programming model, as well as concepts, terminology and abbreviations. Volumes 2...7 present the APIs for the availability management framework, cluster membership service, checkpoint service, event service, message service and lock service, respectively. AIS B.01.01 offers significant enhancements to the availability management framework. For example, it introduces a new component lifecycle interface, provides interfaces for starting and stopping healthchecks and for component-invoked healthchecks, adopts LDAP naming conventions for AIS objects, replaces opaque types with 64-bit integers and extends the redundancy model descriptions. It also aligns the APIs of the various services so that they have common function signatures and follow a uniform programming style.

Numerous companies in the computer and communications industries have adopted the HPI and AIS standards of the SA Forum. Many of them now provide COTS high availability infrastructure components, systems and services. These companies have registered their products as "SA Forum Registered" on the SA Forum web site at www.saforum.org. A current listing of such products can be found at that site. The SA Forum is currently developing a systems management specification (SMS). The SMS provides a simple network management protocol (SNMP) and web-based interface that supports distributed access, monitoring and control of the HPI and AIS services using management information bases (MIBs) and common information model (CIM) schemas. The SA Forum plans to release the SMS in 2005.

The SA Forum hardware platform interface and application interface specification allow programmers to write software that is portable across different hardware platforms, operating systems and high availability middleware. As open standards for COTS components, these specifications result in lower costs and better service for the end users, fewer integration problems for system providers, shorter development times for hardware and software vendors and reduced costs for the end users. The latest versions of the SA Forum HPI and AIS are freely available for all to download from the SA Forum web site at www.saforum.org/specification/download. These specifications serve as the baseline for future certification testing and for future SA Forum specifications, such as the systems management specification. ■

Product News

SBS: CompactPCI InfiniBand host channel adapter

SBS Technologies introduced the IB4X-CPCI-2A InfiniBand HCA that is engineered to drive the performance of high-speed InfiniBand fabrics with low CPU utilization. This 3U-sized InfiniBand HCA card comes with a 6U front panel; a 3U front panel is available on request. The HCA's two 4X InfiniBand ports can deliver overall throughput close to 400MB/s. Configured with 128MB of DDR memory, the IB4X-CPCI-2A provides the high throughput required for embedded data transport applications and high-performance computing clustering (HPCC). The IB4X-CPCI-2A HCA includes circuitry to detect and power an optional IB4X-OMC media converter device that allows the HCA to use fiber so users can conveniently select copper or fiber media on a port by port basis.

News ID 7436

WinSystems: PC/104 GPS receiver and cellular modem

WinSystems introduced the PCM-GPS, a PC/104 module that adds both, GPS and cell modem support, on a single card. The GPS receiver provides complete position, velocity, and time information while the cell modem supports standards-based multiband GSM/ GPRS and CDMA. This board is targeted for both, fixed and mobile applications that require navigation, tracking, data logging, timing and/or communication functions.

News ID 7406

ACT/Technico: 10/100Tx Ethernet modules on PMC form factor

ACT / Technico announced a new line of 10 / 100Tx Ethernet modules on the PMC form factor. The new Ethernet PMCs were designed to fill a void in the marketplace for PMC-based Ethernet boards and are useful for a variety of real-time and embedded applications, including interfaces to broadband and wireless systems at the edge of the network, among others.

News ID 7219