

Service Availability Forum: Creating Open Specifications for Service Availability

By Jon Kenton – Strategic Marketing Manager
Motorola Embedded Communications Computing

Whether in business or our personal lives, dependence on technology is now more than evident. That dependence goes way beyond the personal devices we use directly, such as laptops, PDAs or phones and into the infrastructure that supports them. We live in a connected world where communications holds everything together. Try to going a day without email or a telephone; some might find even an hour tough.

In such an environment, dependence and availability go hand in hand. Based on the literal definition, “dependence” equates to a need for something or somebody to be available in order to exist or survive. Thus availability (defined as the condition of being available) is the fundamental requirement to fulfill the needs of dependence. Reliability used to be a key metric of performance; availability goes beyond how reliable something may be and is a measure of a system or services readiness to perform its function when it is needed, no matter its level of reliability. As we know, reliability in real life situations can not be 100%, therefore highly available systems must account for less than perfect reliability using advanced methodologies, technologies and architectures.

Telecommunications companies have been using Service Availability™ metrics for many years to quantify their equipment’s capability. 5NINES or 99.999% availability is the minimum standard for the majority of networks. Being available for service 99.999% of the time in a 24x7 environment equates to approximately 5 minutes of downtime per year. Downtime includes any time out of service, planned (e.g. HW/SW upgrades) or unplanned. 6NINES capable systems are now being required, which begin to approach 100% availability as the allowed downtime is a mere 32 seconds!

Such levels of availability have been achievable only through the use of highly proprietary and customized architectures and product features, developed to meet the needs of the individual equipment provider and application requirements. Such development and customization is costly and takes many people to maintain. Standardization of these architectures and making products available “off the shelf” will reduce this burden making it easier for applications of all types to meet stringent availability models. The Service Availability™ Forum (SA Forum) is driving this standardization effort to help enable a new open world for Service Availability™.

The Service Availability™ Forum is a standards body comprised of a consortium of industry-leading communications and computing companies. Together these companies support and participate in numerous working groups to develop and publish high availability and management software interface specifications. The SA Forum then works to promote these specifications and facilitate their adoption by the industry as a whole.

The SA Forum mission is to:

“Foster an ecosystem that enables the use of commercial off-the-shelf building blocks in the creation of high availability network infrastructure products, systems and services.”

Service Availability™ Forum Membership

The Membership of the SA Forum covers a broad spectrum, from small software companies to telecommunications giants (see table 1). Members represent the future providers and users of Service Availability™ solutions based on the specifications developed by the forum. All of the companies who are committing resources to this effort recognize the benefits they can glean as well as those that can be realized by the industry as a whole. The SA Forum clearly defines the benefits of adoption at multiple levels from technology creation to service delivery.

Network Service Providers/Operators

- Faster time to market and revenue
- Simplify deployment of new products
- Accelerate innovation
- Equipment provider choice
- Reduce lifecycle costs

Network Equipment Providers

- Increased focus on system definition, integration and value-add
- Faster time to market/revenue
- Broader technology provider choice, pick from “best in breed”
- Reduce product lifecycle costs

Hardware and Software Vendors

- Larger market for standard products
- Reduced need for product customization
- Simplified ecosystems for partnering and more integrated solutions

| | | |
|-----------------------------|-----------------------|--------------------------------|
| ● Artesyn Technologies | ● Intel | ● OSA Technologies |
| ● Augmentix Corporation | ● Kontron | ● Phoenix Technologies |
| ● Clovis Solutions | ● Lucent Technologies | ● Radisys |
| ● Continuous Computing | ● MontaVista Software | ● Siemens |
| ● Ericsson | ● Motorola | ● Solid Information Technology |
| ● Fujitsu Limited | ● MySQL AB | ● Sun Microsystems |
| ● Fujitsu Siemens Computers | ● NEC | ● TietoEnator |
| ● GNP | ● Nokia | ● UXComm |
| ● GoAhead Software | ● Nortel Networks | ● Veritas Software |
| ● Hewlett-Packard | ● NTT | ● Wind River Systems |
| ● IBM | ● Oracle Corporation | ● |

Table 1.

The Specifications of the Service Availability™ Forum


There are three primary SA Forum specifications:

The Application Interface Specification (AIS) defines the interface between the applications and the high-availability (HA) middleware. By using the features of the AIS the application can remain independent of the underlying platform. This greatly enhances portability and enables a very robust HA stack management structure. The fundamental concept on which AIS builds is that of clustering. By taking advantage of redundant hardware and software components the AIS is able to increase Service Availability™ through masking the four possible failure zones: the application, the middleware, the operating system and the hardware. The AIS itself is a C language API that provides a number of services to manage a distributed mechanism for supporting cluster membership, application failover, and checkpointing. Through the creation of multiple nodes (even within a single system) Service Availability™ is improved by distributing the application and middleware across the logical and/or physical nodes.

The Hardware Platform Interface (HPI) defines the interface between the hardware and the HA middleware and makes each independent of the other. As with AIS the HPI enables portability, in this case it allows for platform vendor independence. It also means that an application can be created that can span multiple hardware architectures while remaining compatible from a management perspective. Every hardware platform has similar elements, although they can differ significantly in implementation. The HPI creates an abstracted view of the platform hardware-specific characteristics of a system and provides a standardized methodology for monitoring and control. The abstracted view is divided into “entities” each of which have associated controls, sensors, entity inventory repositories, and watchdog timers. These elements provide the HPI with the means to manage the overall state and health of the underlying hardware platform.

The Systems Management Specification (SMS) is the newest of the specifications developed by the SA Forum and will be fully released later this year. The SMS defines the interfaces to access the monitoring and control aspects of the AIS and the HPI interfaces, as well as a comprehensive notification interface for HA systems. The overall goal of the SMS is to address the administrative operations and management of cluster configuration, platform system model configuration, deployment and statistical data monitoring.

Figure 1 shows how the three specifications will work together in a layered architecture that builds from the foundation of the hardware platform to the end application that resides at the top of the stack.



Service Availability™ Forum Resources

There are various resources available through the SA Forum website www.saforum.org.

The specifications themselves may be downloaded directly from the site. Each of the specifications available contain an overview description of the specification, services, dependencies and various usage models, API specific definitions, naming conventions and policies.

Hardware Platform Interface (HPI) specification: The full specification and header files for the HPI are available for download, either together or separately

Application Interface specification (AIS): The full specification and header files for the AIS are available for download, either together or separately

Systems Management specification (SMS): The interface specification that enables the service event and error reporting for AIS and HPI. The SMS will be available for download in mid 2005

The product registry and developers forum are additional resources available through the website.

The Service Availability™ Forum Product Registry: The registry is a program that enables members to post their implementation methodology to the Service Availability™ Forum Product Registry for evaluation and testing. Compliance with all registration requirements enables members to label their products as "Service Availability™ Forum Registered." As more offerings become available from the community, the registry will act as a guide for users and adopters of SA Forum compliant products to help select those that best match their specific requirements.

SAF-HPI Developer's Forum: This online facility allows developers to submit inquiries to the 'HPI Developers' Forum. The idea is to obtain feedback on implementation issues from others who have already been through the process.

Consortium Publications

The SA Forum website also provides access to various documents & white papers including:

- What is the Service Availability™ Solution?
- The Service Availability™ Forum Platform Interface
- Implementing HA databases within an SA Forum AIS-compliant Framework
- SA Forum datasheet: provides a brief overview of the goals, objectives and deliverables of the Forum

Working Groups

The engine room of the SA Forum is powered by three primary working groups: The Technical, Marketing and Compliance Workgroups. It is within these groups that the contributions from each of the member companies come together to create the specifications.

Technical Workgroups: These very active groups represent the core of the forum. Each subgroup is responsible for a distinct area and their collective output combine to create the SA Forum specifications.

- Application Interface Specification Subgroup (123 members)
- Application Services Subgroup (51 members)
- Architecture Subgroup (70 members)
- Hardware Platform Subgroup (64 members)
- SAF-OSDL Subgroup (18 members)
- System Management Subgroup (111 members)

Marketing Workgroup (62 members): It is the role of the MWG to promote adoption of SA Forum specifications. This is achieved through an active recruitment campaign and conventional PR along with outbound marketing activities. The MWG is also very much involved in establishing requirements and working with the technical subgroups to ensure that delivered specifications meet the needs of the market.

Compliance Workgroup (24 members): As with any standardization effort, ensuring adherence and compliance to the spec is crucial to the goals of portability and compatibility. At this time the SA Forum maintains a product registry where members can submit self-certified documentation related to the conformance of their specific offerings to the relevant specification(s). The Forum is working on an independent compliance mechanism through a third party. This will allow any vendor to submit products for testing, the results of the tests can be used to verify compliance to each of the SA Forum specifications. Full details of the compliance testing program will be announced later this year.

Working Group Collaboration

Participation and collaboration within any standards body represents similar challenges. Members of the various working groups are from different companies and often different countries. This makes electronic and asynchronous communication the essential medium if any work and significant output is to be achieved. Email, website document repositories and message reflectors are the primary methods used by all the groups for ensuring everybody in the work group stays in lockstep. Most groups hold frequent (usually weekly) conference calls where topics and issues can be discussed live. The Forum holds quarterly “face-2-face” meetings that the majority of workgroup members attend. Lasting two or three days these meetings allow for concentrated and detailed sessions to be held with each workgroup. They are usually very lively and productive meetings

and set the tone for the next period of electronic activity. While such a modus operandi may sound difficult, it allows members to balance their “day jobs” with the needs of the Forum. Moreover it ensures the goals and objectives are monitored and achieved on schedule.

Overall, the benefits of membership to the companies and individuals involved are numerous. As discussed earlier, standardization brings technological and economy of scale benefits to the industry through the collaboration and interaction of industry experts in their respective fields. All members are able to realize the opportunity to help develop and shape the future of highly available systems solutions while networking with their peers.

The use of commercial off-the-shelf system platforms in the communication infrastructure is inevitable. Developing applications that can leverage this equipment, while maintaining the availability needs of carrier-grade systems, is a challenging task. However, by creating Service Availability™ interface specifications, and deploying building blocks that conform to those specifications, the members of the Service Availability™ Forum are creating a new technology ecosystem for the rapid development of solutions which can meet that challenge.