Introduction to Service Availability Forum

Sasu Tarkoma (sasu.tarkoma@cs.helsinki.fi)

Seminar on High Availability and Timeliness in Linux
University of Helsinki, Department of Computer Science
Spring 2003

Abstract

The Service Availability Forum is a coalition of communication and computing companies advancing open and standard interface specifications for carrier-grade service availability. The Forum's service availability solution is to combine high availability and service continuity in order to provide continuous user experience of services. This paper presents an overview of the forum and examines the two specifications that the forum is developing: the Application Interface specification and the Hardware Platform Interface (HPI) specification.

1. Introduction

The Service Availability Forum (SAF) is a coalition of communication and computing companies advancing open and standard interface specifications for service availability [SAF03]. The companies include Nokia, IBM, Intel, MontaVista and other key players of the industry. The key motivation for this forum is the transition to packet-based multi-service networks that requires a carrier-grade infrastructure. Packet-based transport enables new flexible, scalable and cost-efficient services. Standard legacy Public

Switched Telephone Network (PSTN) technology requires significant amount of data processing to set up or tear down a call; however, failure of the computing components have generally had little impact on calls in process. The new packet-based services change this scenario by relying more on computing components in provisioning the service (audio, data,..), and while the cost is reduced the reliability and availability of these components needs to be ensured. Users expect on demand service delivery without interruptions [SAF02a] [SAF02c].

Different "carrier grade" systems that are sufficiently dependable have been built with off-the-shelf resources. They are generally called as COTS, commercial off-the-shelf building blocks. These systems have included manageable redundancy, and they are available from different vendors; however, different products have different interfaces for managing redundancy and availability. The goal of the Service Availability Forum is to improve this situation by developing the missing standard interfaces [SAF02a].

The high availability infrastructure is based on interoperable hardware and software building blocks, management middleware and applications. The open specification supports interoperability between the implementations of different vendors. Interoperability allows the delivery of continuous voice, data and multimedia service to carriers and end-customers [SAF03].

The SAF will develop two carrier-grade interface specifications that will be operating system (OS) and platform agnostic. The two initial specifications are the Application Interface (SAI-A), and the Platform Interface (SAI-P) specifications. This paper is structured as follows: Section 2 presents the structure and timeline of the Service Availability Forum. Section 3 examines the specifications, and Section 4 presents the conclusions.

2. Service Availability

2.1 Overview

The Service Availability Forum is using a two-step process in order to standardize the interfaces:

1. Analysis by industry experts of what is needed from a platform in order to support carrier-grade dependability, and
2. the development of standard interfaces that can be implemented by platform and middleware developers.

The key features being standardized by the Service Availability Forum are:

- High availability
- Service Continuity

The aims of the Forum are to eliminate lost connections, late voicemail delivery, slow data transmissions, and reduce cost and improve development times of the design and manufacturing of network components and services. The interfaces are intended to be used on the following types of equipment: circuit-switched telephony and data network, cable-telephony network, wireless and Internet network, and other packet-switched networks [SAF02a].

2.2 Background

The Forum was announced in December 2001 as the successor to the HA (High Availability) Forum that pioneered the standard-based architecture high availability systems. HA produced a whitepaper titled "Providing Open Architecture High Availability Solutions". Two basic requirements were identified [HAF01]:

- High availability: 99.999% system uptime (or better) for on-demand services.
- Service continuity: customer sessions and state are maintained for uninterrupted services.

The HA whitepaper also presents customer requirements for open HA systems, configuration management, and layer specific methods and capabilities that are required. In addition, the SAF Hardware Platform Interface (HPI) draws heavily on the Intelligent Platform Management Interface (IPMI) specification that defines platform-independent data formats and capabilities [SAF02b].

2.3 High Availability

The starting point for the SA Forum has been high availability, which ensures that systems are up and running, and the probability of a breakdown and the expected downtime are under the specified limits. Availability is the measure of the probability that a service is available at any given instant. Currently, systems that have an uptime of five-nines (99.999%) are called highly available. These systems have less than five minutes of cumulative downtime per year. Commonly PSTN provide 6nines availability. Availability is a function of reliability, reparability, and supported by redundancy. The reliability is usually expressed as a Mean Time Between Failures (MTBF) on a per component-basis. Reparability tells us how quickly the system can be repaired after the occurrence of failures, and expressed as a Mean Time to Repair (MTTR). In the absence of any redundancy, the availability is given by the formula, *Availability = MTBF/(MTBF+MTTR)* [SAF02a] [HAF01].

A key requirement for carrier-grade availability is manageable redundancy. In this context, redundancy means that when one component in the system fails, a duplicate or a replica component will be available and can continue processing. Manageable redundancy requires that the platform can detect failures, and automatically reconfigure itself. Redundancy adds to the reparability by having a backup or standby for the components of the system. Redundancy helps in driving the system MTTR towards zero. Key measures for redundancies are the amount of redundancy, and how quickly backup components can be activated.

2.4 Service Continuity

The High Availability Forum sees high availability as a basic requirement for continuous computing; however, if a component fails and a redundant copy is used, what happens to user experience and user sessions? The system needs to ensure that not only hardware, but also the service is delivered uninterrupted. Therefore, we need a continuity solution to preserve customer data, and session state across fault scenarios.

Service Continuity means that end-user sessions are maintained even if some individual systems or components fail. For example, the state of the application sessions for each client must be maintained during switchover scenarios.

2.5 Standardization Process and Timeline

The Forum works through collaborative process, where members participate in request for proposals (RFP), and all submissions undergo evaluation and a voting process. The iterated responses will then be compiled into one application interface baseline document and one platform interface baseline document. These documents are the starting points for the interface specifications. The Forum organizes face-to-face meetings approximately every 2 months.

The Forum released a request for proposals for programming interface specification between middleware and the platform on January 2002. The platform specification was released on 10.07.2002 for public review and is available for download. The Forum plans to release the Application Interface specifications early this year. Potential implementers and usage licensors will use a web-based template to disclose information on compliance and testing. The Forum intends to develop a formal process for compliance certification during 2003.

3. Specifications

Initially, the Service Availability Forum will develop two interface specifications that are operating system, and platform agnostic. The two initial specifications are:

- The Service Availability Forum Application Interface (SAI-A) API between applications and SA middleware.
- The Service Availability Forum Platform Interface (SAI-P) interface between SA middleware and platform components [SAF02b].

Figure 1 illustrates the different layers in the service availability architecture. The interfaces emphasize communications between components, and distributed processing with redundancy. The Forum is implicitly defining a software layer called Service Availability Middleware, which is carrier-grade software that conforms to the interface specifications. This middleware needs to have management support for redundancy, fault detection and reaction, and the basic tools that the applications need to respond to system fault conditions. The interfaces abstract the details of the underlying system platform and allow applications to continue without loss of service continuity.

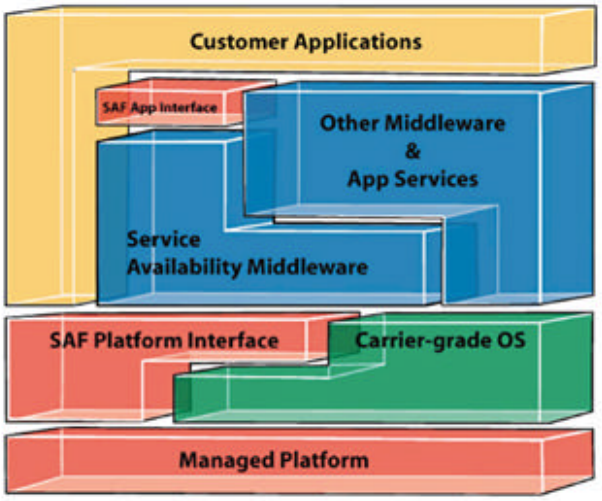

Figure 1 The Service Availability Interfaces (Source [Hol02])

3.1 Application Interface

The SAF Application Interface defines the services (API) that the SA middleware needs to provide for applications. The interface provides access to tools that support the distribution of application software processing over multiple computing elements. In addition, the interface allows applications to respond to failures of those elements without loss of service. Middleware that conforms to this interface is called Service Availability middleware (figure 1).

In the Service Availability model, the middleware supports monitoring and controlling of the various physical components of the platform using the Platform Interface. This includes support for "hot-swappable" components and also changing the hardware configuration. The middleware provides services for applications independent of the hardware.

3.2 Platform Interface

The Hardware Platform Interface (HPI) is an open interface for accessing the operating system and hardware components of the computing platform. The interface is platform and operating system agnostic, and its primary user is the Service Availability middleware software. The platform interface is an independent specification, and any software package may use it that needs to monitor and control the hardware platform [SAF02b].

HPI provides a common model of the hardware platform and its management capabilities. The components of the system are represented as entities. Each entity has a unique identifier, which is called an entity path in the specification. The entity path is defined by the component's location in the physical hierarchy of the system, for example: "Power supply number 2, contained in Subrack number 4, contained in Rack number 3." An entity has an associated resource data record, which stores information pertaining to manageability, which is modeled using controls, watchdog timers and sensors. HPI users

receive information about the state of the system by using these mechanisms. Entity management may include any combination of the following functions:

- Reading values pertaining to component status (modeled with "Sensors").
- Controlling the operation of a component (modeled via "Controls").
- Reporting inventory and static configuration data (reported via "Entity Inventory Records").
- Placing watchdog timers on components (modeled via "Watchdog Timers").

Entities are grouped under resources. Resources support runtime insertion, removal, monitoring and controlling of components. Runtime changes are reported as "Hot Swap" events, which the resources store in event logs for later retrieval. Resources are divided into one or more domains, where a domain provides access to some set of resources within the system. Domains provide a way to partition and separate tasks, and they may overlap. Events can be subscribed from different domains. A domain controller provides information about the resources in a domain, maintains a Resource Presence Table (RPT), and manages events. All access to the HPI is via a session, which is open to one domain at a time; however one HPI user may have multiple sessions open at once. The user of the HPI is an operating system, middleware etc. Figure 2 presents a high-level overview of the HPI domain.

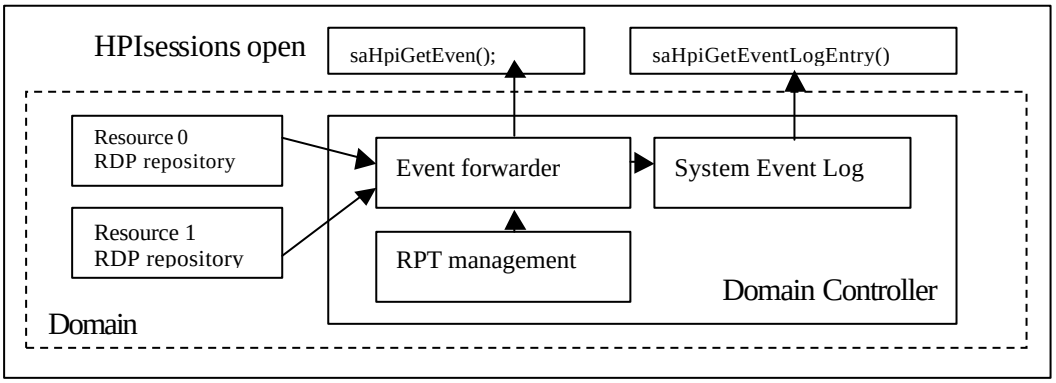

Figure 2 HPI domain with two open sessions [SAF02b]

The platform interface will be made available by vendors as a set of C-language library calls, and a header file from the SAF specification. The vendors modify the header file to

use correct basic data types for the processor family. The library functions are written by the vendor to map to the HPI model and to provide the actual functionality to the hardware platform capabilities.

3.3 Industry Support

Companies will use implementations of the specifications to setup, monitor and perform fault-recovery of different network equipment solutions. The Forum has worked closely with PICMG (The PCI Industrial Computer Manufacturers Group) and other standards organizations. The Carrier Grade Linux Platform Working Group intends to align with existing industry forums, such as SAF [OSD02a]. Figure 3 presents the role of SA Forum interfaces in the OSDL's Carrier-Grade Linux Platform Architecture**.** OSDL will ensure that the Carrier Grade Linux provides suitable interfaces for supporting highly available middleware. CGL will be compliant with the SA Forum standards where there is specific overlap [OSD02].

MontaVista plans to use its IPMI (Intelligent Platform Management Interface), a standard for monitoring the hardware in a system, driver and libraries to support the SA Forum's Platform Interface API during 2003.

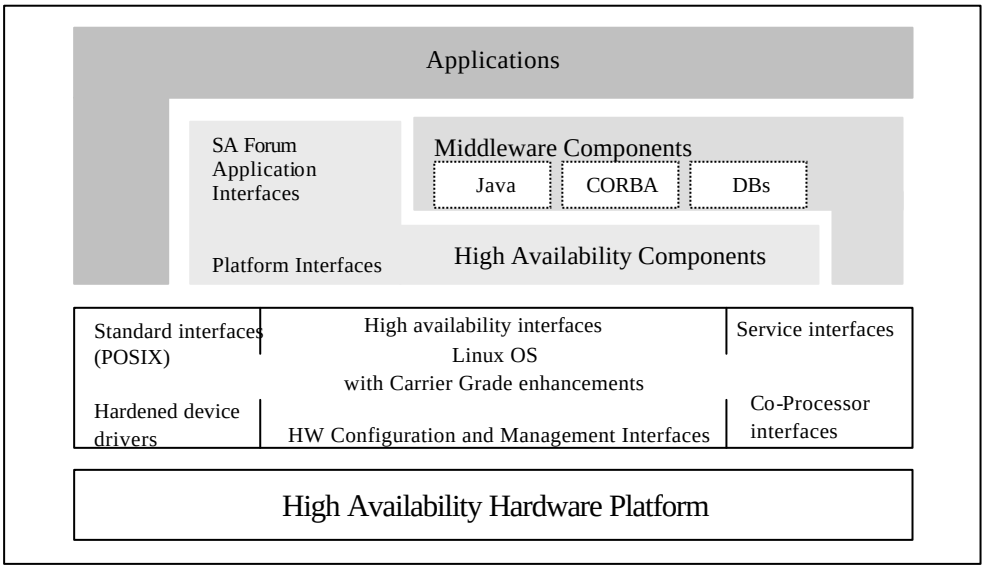

Figure 3 The SAF interfaces in Carrier-Grade Linux Platform Architecture [OSD02a]

4. Conclusions

The Service Availability Forum will develop two carrier-grade interface specifications that will be operating system and platform agnostic. The key features being standardized by the Service Availability Forum are: high availability and service continuity. Through the specifications the Forum is defining a carrier-grade software layer called Service Availability Middleware. This middleware needs to have management support for redundancy, fault detection and reaction, and the basic tools that the applications need to respond to system fault conditions. The SAF Application Interface provides access to tools that support service continuity and the distribution of application software processing over multiple computing elements. The Platform Interface specification addresses issues pertaining to monitoring and controlling, and defines the interfaces between the middleware, and the operating system and hardware. The interfaces abstract the details of the underlying system. The Forum has released the hardware Platform Interface specification for public review and plans to release the Application Interface specification early this year.

References

[HAF01] High Availability Forum. Providing Open Architecture High Availability Solutions. Revision 1.0, 2001. Available at: ftp://download.intel.com/platforms/applied/eiacomm/papers/ha_solutions.pdf

[Hol02] Colin Holland. Service Availability Forum's Platform Interface spec released to ease infrastructure design. Embedded Systems (Europe), November 2002. Available at: http://www.embedded.com/story/OEG20021108S0021

[JP02] Timo Jokiaho and Dave Penkler. The Service Availability Forum Platform Interface. RTC Magazine, December 2002. Volume XI, Number 12. Available at: http://www.rtcmagazine.com/pdfs/2002/12/rtc12_techfocus.pdf

[OSD02] OSDL. Carrier Grade Linux Technical Scope. White Paper, June 2002. Available at: http://www.osdl.org/projects/cgl/cgl_technical_scope_whitepaper.pdf

[OSD02a] OSDL. Making the Case for Carrier Grade Linux Platform, 2002. Available at: http://www.osdl.org/pdf/case_for_carrier_grade_linux.pdf

[SAF03] Service Availability Forum website, 2003. Available at: http://www.saforum.org

[SAF02a] Service Availability Forum. Whitepaper, 2002. Available at: http://www.saforum.org/specification/specification_white_paper.pdf

[SAF02b] Service Availability Forum. Hardware Platform Interface Specification. 2002. Available at: http://www.saforum.org

[SAF02c] Service Availability Forum. Standards for a Service Availability Solution. 2002. Available at: http://www.saforum.org/about/solution_backgrounder.pdf