

Developing High Availability xTCA Applications with SA Forum Based Middleware

John Fryer

SA Forum Marketing Work Group Chair –
Service Availability Forum

Director, Software Business Development –
Emerson Network Power

Defining Service Availability

- Service availability is an extension of high availability that implies a service is always available regardless of hardware, software or user fault.
- High availability is a system design and associated implementation that ensures a certain absolute degree of operational continuity in a given measurement period.*
- Key principles of service availability:
 - Redundancy – “backup” capability in case of need to failover due to a fault
 - Stateful and seamless recovery from failures
 - Mean time to repair (MTTR) – time to restore service after an outage
 - Fault prediction & avoidance – take action before something fails

*Source - Wikipedia

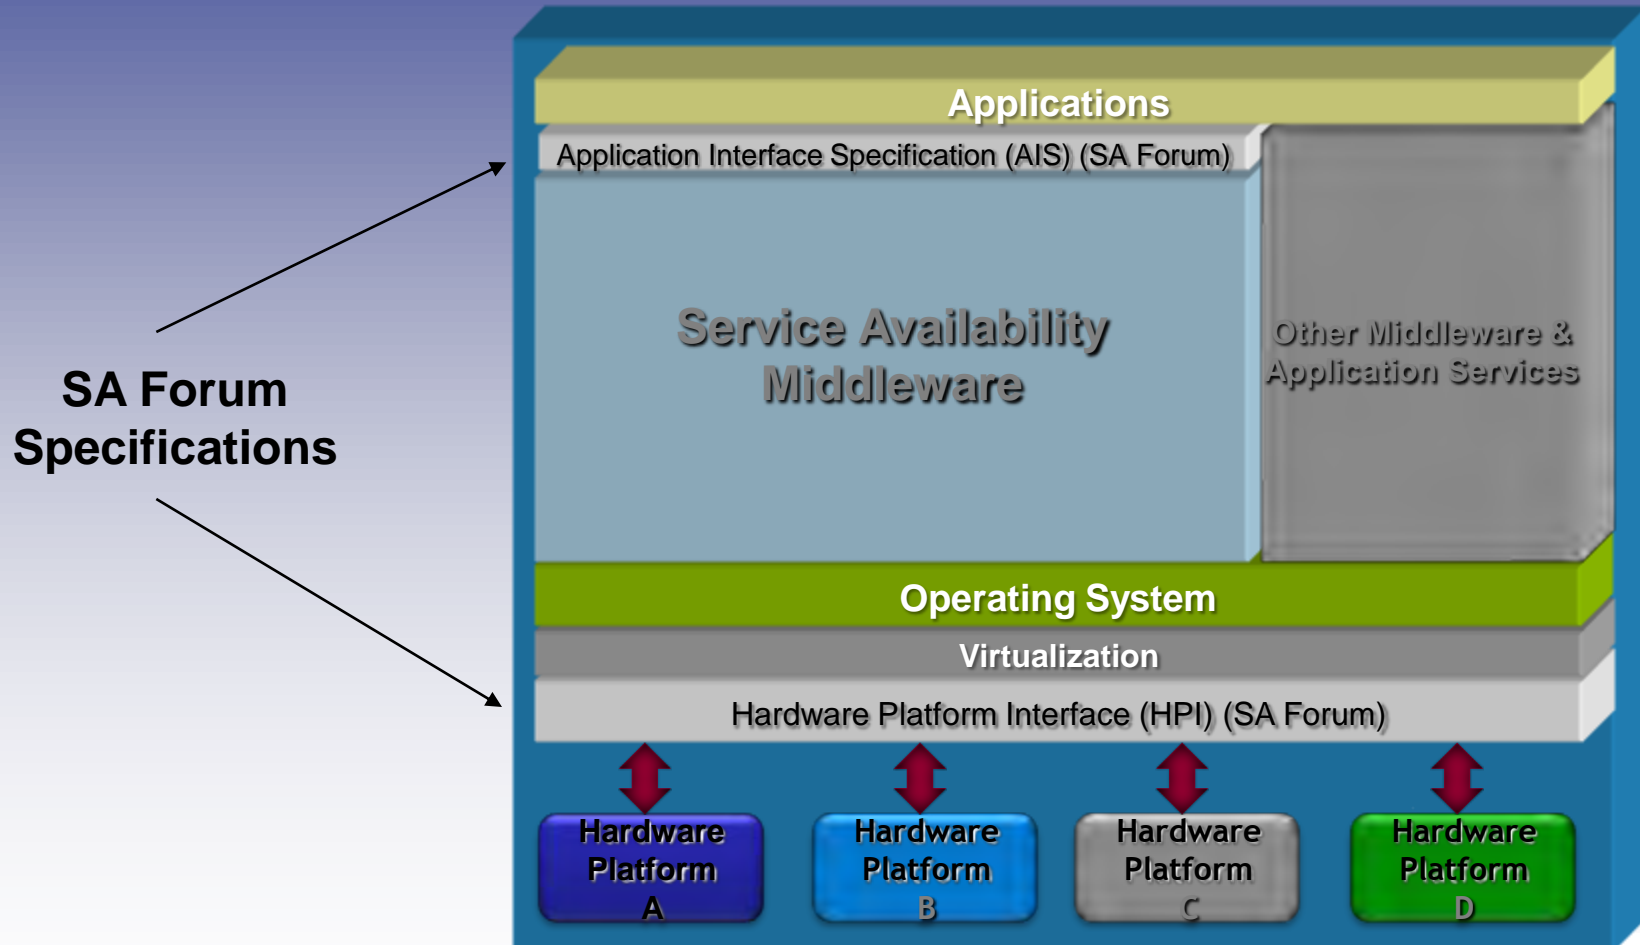
Redundancy – Part of the Solution

- Problem
 - A real system has a large number of hardware and software components
 - Hard to achieve high availability by increasing the mean time between failures (MTBF)

- Solution
 - Use redundant components to achieve high availability by decreasing the MTTR to close to 0
 - From hours or minutes to a few milliseconds
 - AdvancedTCA dual star architecture provides hardware aspects

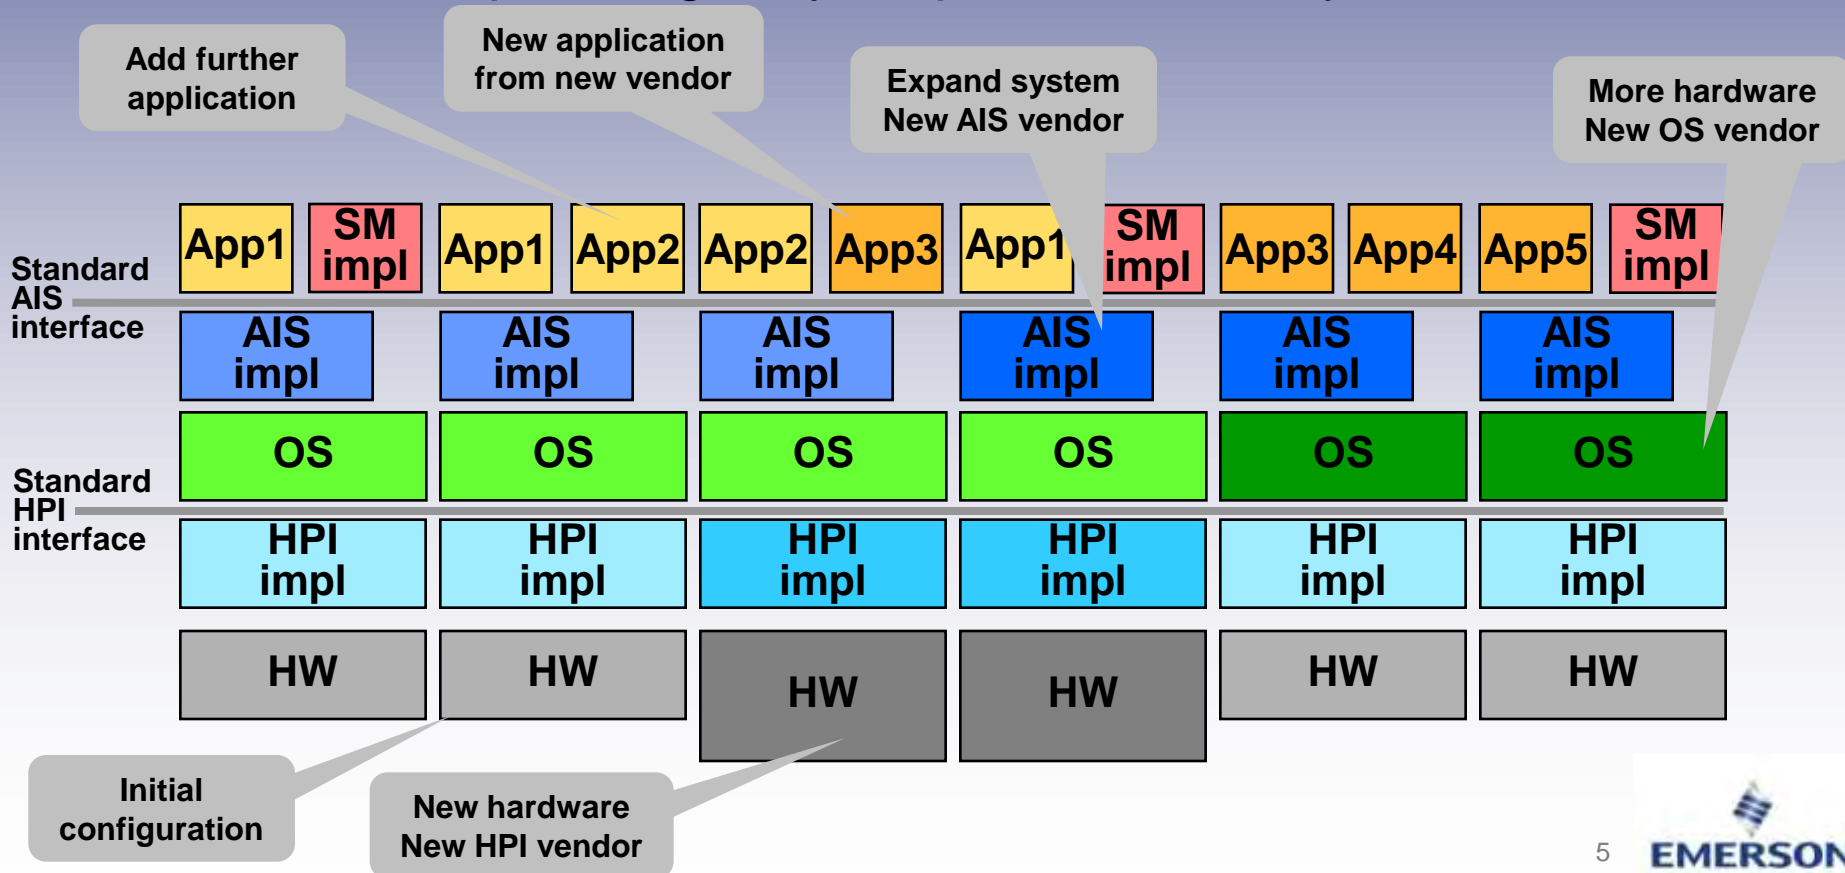
**A Homogeneous System Approach to Integrate
Hardware and Software
Middleware is a Key Requirement**

Service Availability Reference Architecture



Architecture Application to ATCA

- AdvancedTCA is a loosely coupled system of compute and I/O nodes
- SA Forum specs logically map to an ATCA system

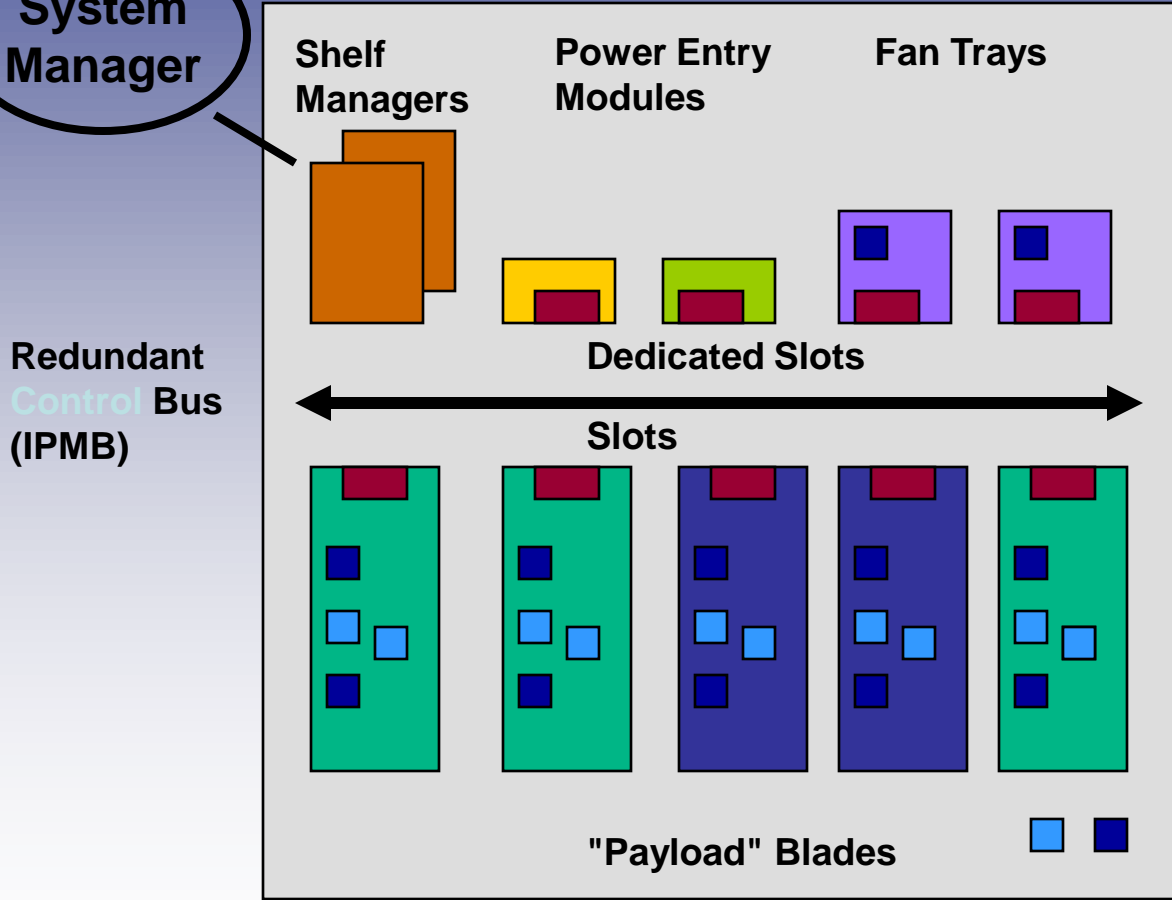


HPI-to-AdvancedTCA Mapping

- The PICMG AdvancedTCA specification is perceived to be an underlying hardware platform for SA Forum Hardware Platform Interface (HPI) implementations
- The SA Forum HPI presents a data model and programmatic interface to describe generic hardware platforms
- The HPI-to-AdvancedTCA Mapping specification details how to map the HPI to the AdvancedTCA specification
 - It exposes AdvancedTCA Shelf Management functionality and data in a standard, vendor-independent manner via the HPI
- The AdvancedTCA Shelf Manager capabilities may be used by an application that provides overall system management
 - The HPI provides the interface between the AdvancedTCA Shelf Manager and the HPI user in an abstract, easy-to-use manner

Example AdvancedTCA Shelf

**System
 Manager**



← **IPMI Controllers**

← **Various sensors and controls**

Shelf

Application Interface Specification

- The Application Interface Specification (AIS) is a set of open standard interface specifications
- The AIS defines an Application Programming Interface (API) for middleware between the applications and the operating system
- The AIS is divided into the following areas:
 - Availability Management Framework (AMF)
 - Cluster Membership Service
 - Checkpoint Service
 - Event Service
 - Message Service
 - Lock Service
 - Notification Service
 - Log Service
 - Information Model Management Service

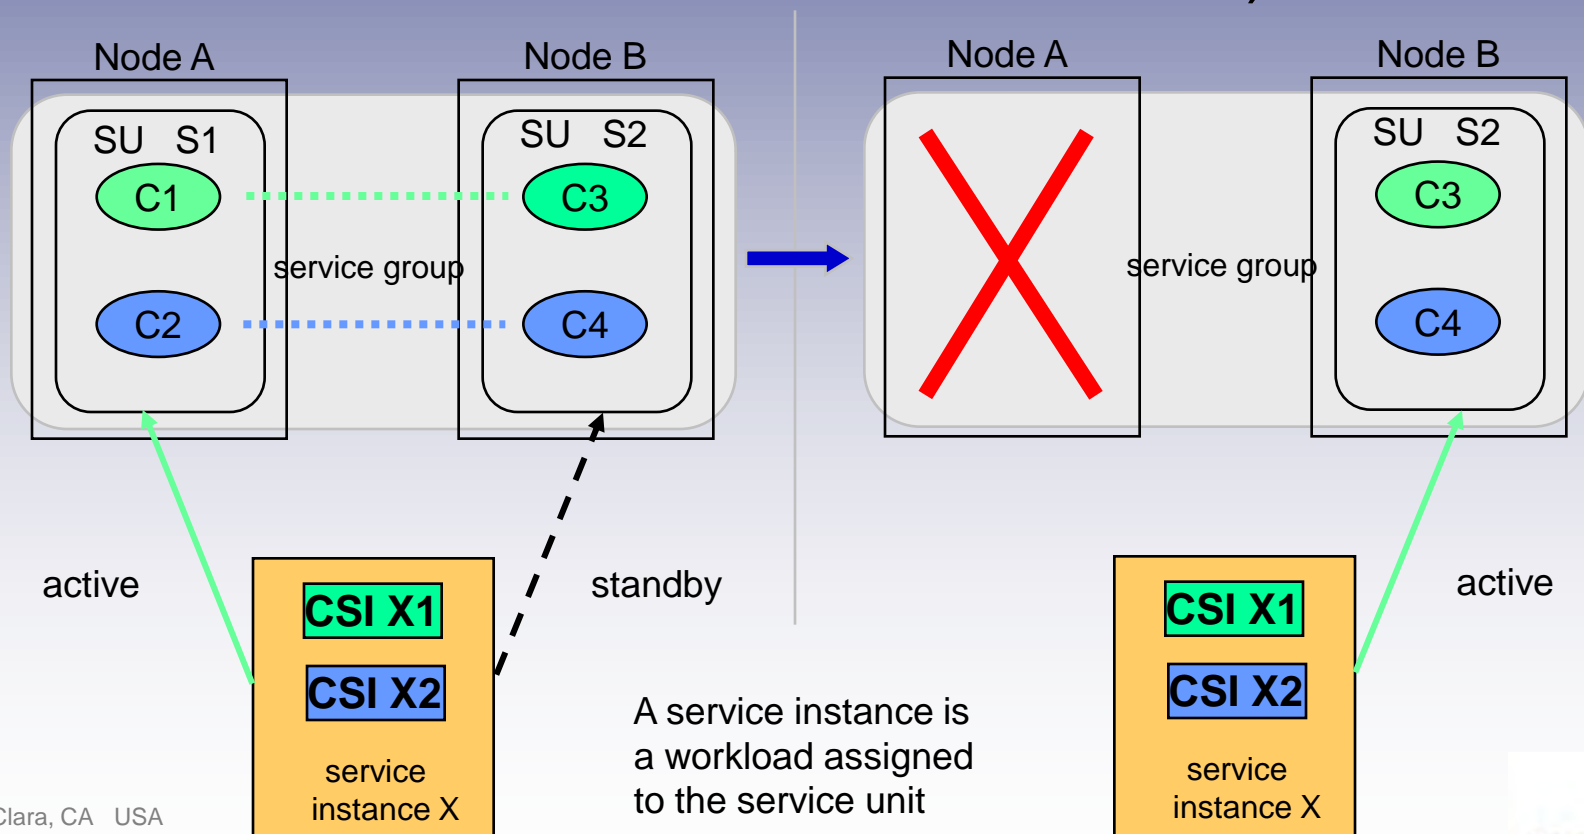
ATCA and AIS – Service Availability

- AMF – define mechanisms and actions to relate failures
 - Hardware and software
 - Concept of physical and logical resources
 - Failover, restart, ignore etc.
 - Enables an integrated system approach
- Event Service
 - “Service” affecting actions into AMF
 - Hardware reporting mechanism via HPI

Instrumenting ATCA Hardware with Appropriate Sensors and Controls is Essential to Achieve a Service Availability System

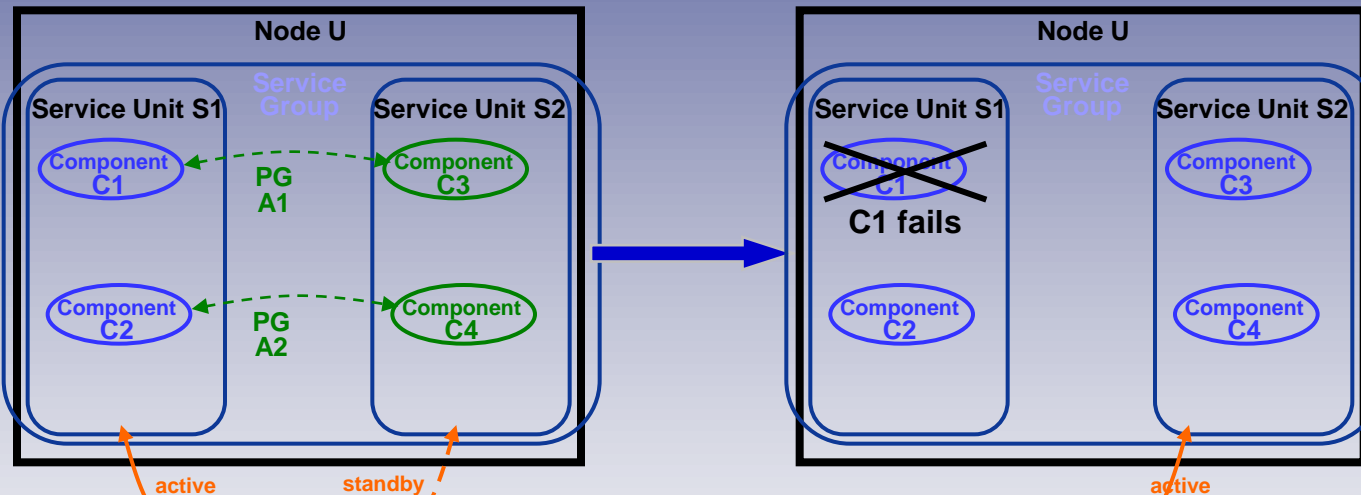
Application Interface Specification

- Example of the 2N redundancy model (with two service units on different nodes)



Potential Multi-Core Usage

- Example: Two service units on the same node and a component fails



■ After a fault that disables component C1 in service unit S1, service unit S2 is assigned to be active for service instance A

■ Note that a fault affecting any component in a service unit, that cannot be recovered by restarting the affected component, causes the entire service unit and all components within the service unit to be withdrawn from service

■ Here, even though component C2 is still fully operational, it must failover to component C4

SA Forum Extended Training Material

- The SA Forum Extended Training Material comprises:
 - Service Availability Interface Specifications
 - Overview
 - 2 lessons on the Service Availability Interface
 - Hardware Platform Interface
 - Overview of the HPI
 - 4 lessons on the HPI
 - 2 lessons on the HPI-to-AdvancedTCA Mapping
 - Application Interface Specification
 - Overview of the AIS
 - 4 lessons on the Availability Management Framework
 - 1 lesson on each of Cluster Membership Service, Checkpoint Service, Event Service, Message Service, Lock Service, Log Service, Information Model Management Service
 - 2 lessons on the Notification Service
 - System Management Interface
 - 1 lesson on HPI-SNMP Distributed Systems Management
 - 1 lesson on AIS-SNMP Distributed Systems Management
 - Application Webcast Series
 - Aimed to educate viewers on implementing open specifications for carrier-grade and mission-critical applications and systems

Thank You

For more information, please visit
www.saforum.org

John Fryer
John.Fryer@Emerson.com