

High Availability with the openais project

***Prepared by:
Steven Dake
October 2005***

Agenda

- Service Availability Forum
- Reliability and Availability
- Application Interface Specification
- The openais project

Service Availability Forum – Mission

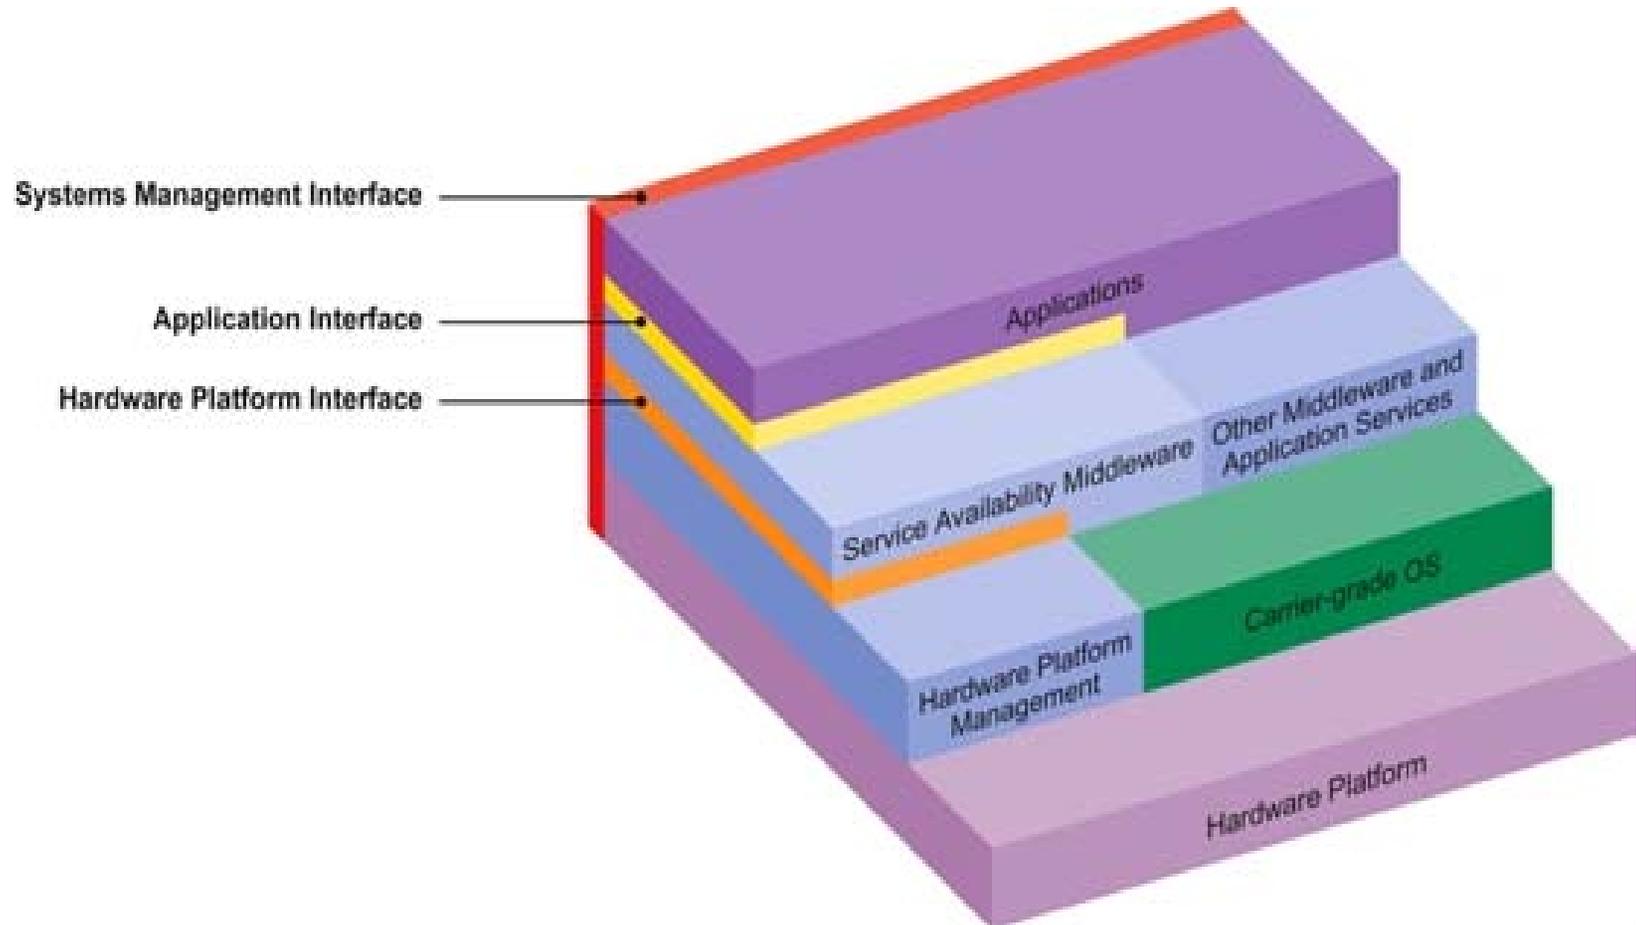
The Service Availability™ Solution helps meet end-user expectations for voice, data and multimedia services delivered with the dependability of traditional telecommunications.

The Service Availability™ Forum is addressing this by fostering an ecosystem to enable the use of commercial off-the-shelf building blocks in the creation of high availability network infrastructure products, systems and services. The Service Availability™ Forum will accomplish this through developing and publishing high availability and management software interface specifications as well promoting and facilitating their adoption by the industry.

Service Availability Forum Member Companies

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

Service Availability Forum – The Software Stack

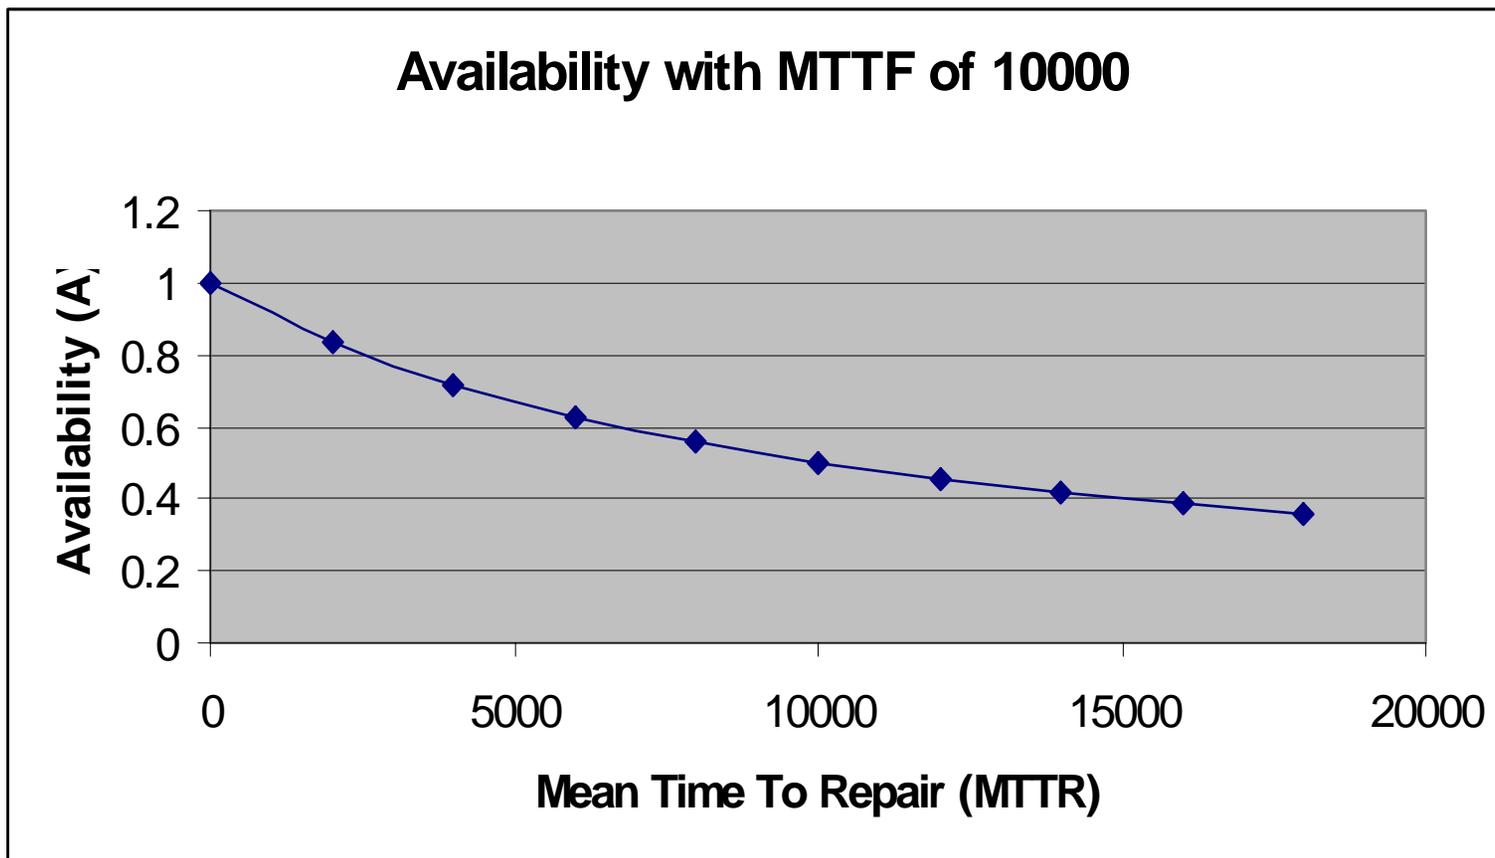


Reliability and Availability – Availability Equation

$$A = \frac{MTTF}{MTTF + MTTR}$$

Where MTTF is the mean time to failure and MTTR is the mean time to repair.

Reliability and Availability – Availability with fixed MTTF and variable MTTR



Application Interface Specification - Overview

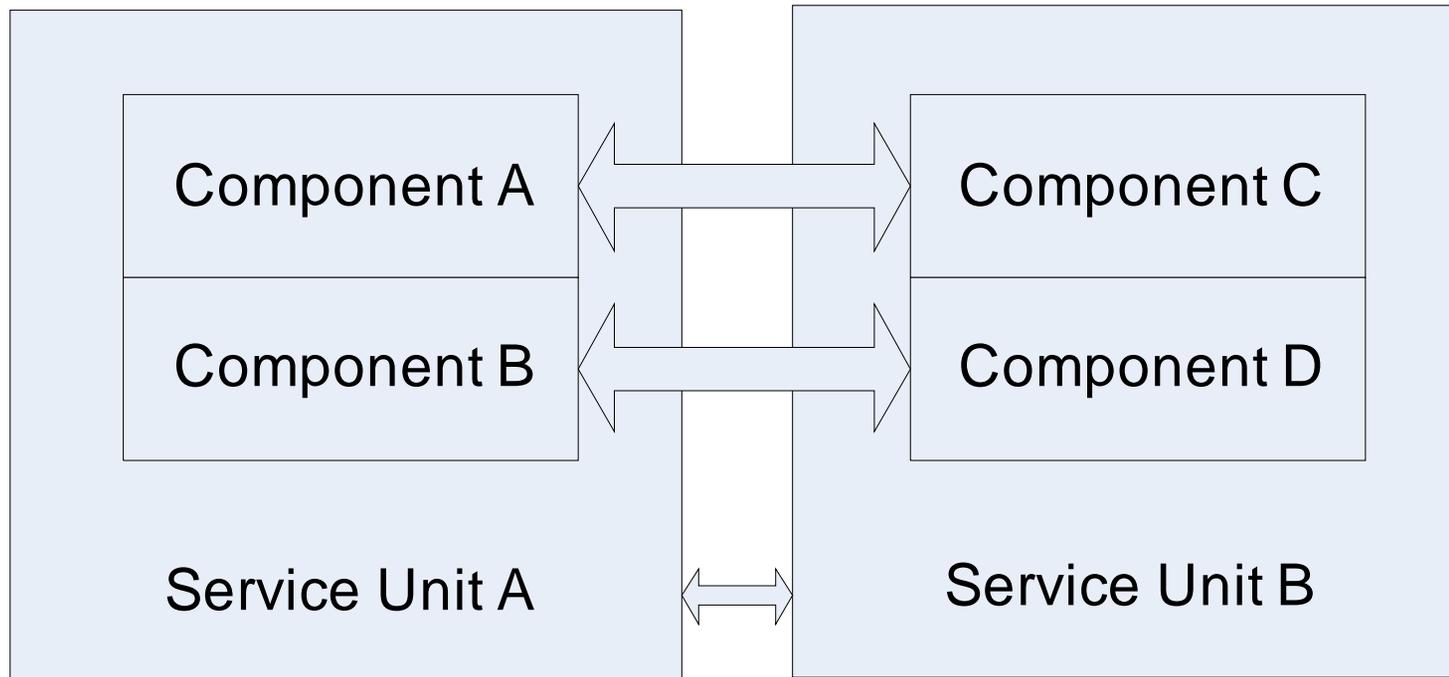
- High Availability Specification
- Application Failover
 - Checkpoint Service
 - Availability Management Framework
- Communication
 - Cluster Membership Service
 - Event Service
 - Message Service
- Mutual Exclusion
 - Distributed Lock Service

Availability Management Framework

- Overview

- Allows service to be registered or unregistered
- Instantiates services as active or standby
- Detects service faults
- Provides mechanisms to gather instantiation state
- Mechanism to enable and disable services
- Allows reporting of errors and canceling errors

Availability Management Framework – Service Group



Checkpoint Service - Overview

- Checkpoints are named
- Checkpoints have sections which store data
- Checkpoint sections can be read and written
- When an standby component is directed active by AMF, standby reads checkpoint sections and recovers state

Cluster Membership Service - Overview

- Maintains view of current configuration
- Allows for asynchronous notification of configuration changes via tracking API
- Provides mechanism to read current configuration

Eventing Service - Overview

- Provides named event channels for publish and subscribe
- Publish events to an event channel
- Callback executed when filtered event is delivered
- Events can be filtered by api

Messaging Service - Overview

- Named queue identifiers for sending and receiving messages
- Mechanism to send a request and wait for the response
- Load balancing messages

Locking Service - Overview

- Resources can be locked and unlocked
- Asynchronous notification of many operational types
- Locks can be reclaimed in case of failure of locker
- deadlock detection

The openais project - Agenda

- Setup and Configuration
- Project History
- Architecture
- Performance
- Project Statistics

openais – setup and configuration

- Create shared key:

```
Linux# ./keygen
```

OpenAIS Authentication key generator.

Gathering 1024 bits for key from /dev/random.

Writing openais key to /etc/ais/authkey.

- Save /etc/ais/network.conf:

```
Bindnetaddr: 192.168.1.0
```

```
Mcastaddr: 226.94.1.1
```

```
Mcastport: 6000
```

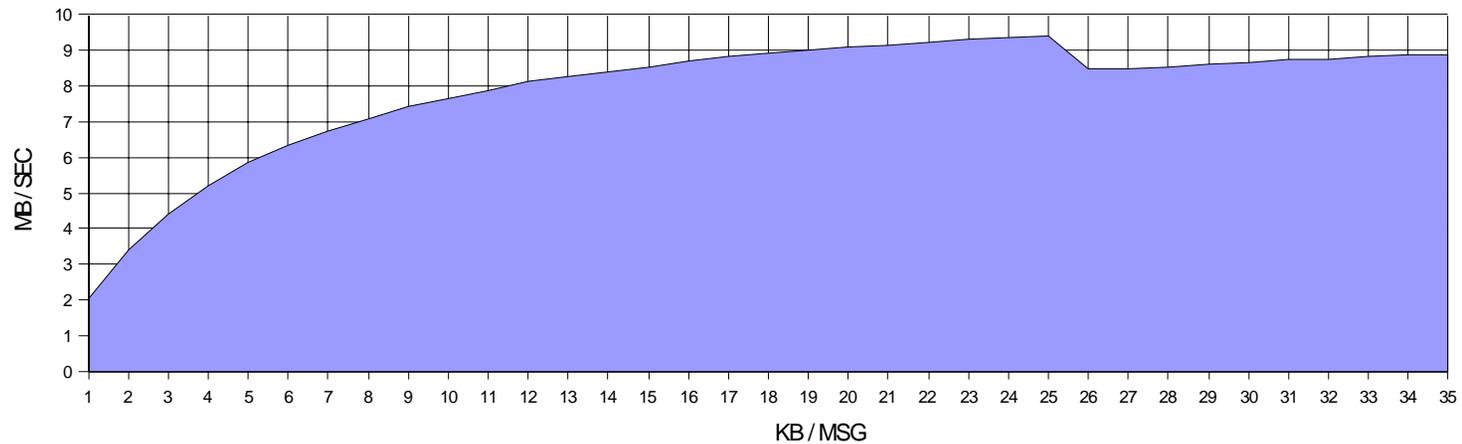
Read QUICKSTART file in source package for more details.

openais – project history

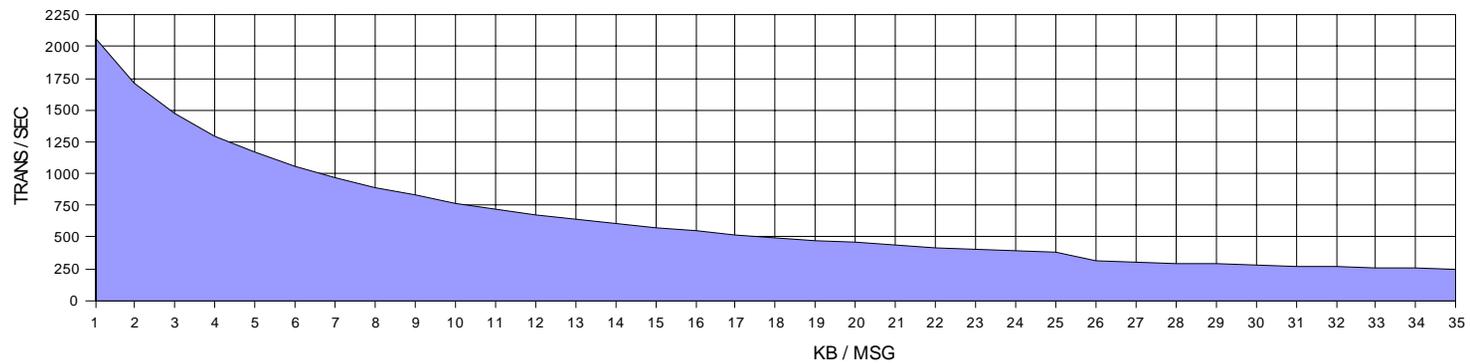
- Project started in January 2002 to support hotswap on ATCA chassis
- Morphed into SA Forum in April 2003
- Virtual Synchrony merged January 2004
- Released to open source under Revised BSD license by MontaVista Software in June 2004 as the openais project hosted at Open Source Development Labs.
- Event service merged September 2004
- Open Source Development Labs and SA Forum officially announce via press release their support for the openais project in November 2004.
- 3rd generation implementation Virtual Synchrony protocol merged January 2005

openais – checkpoint performance

Throughput

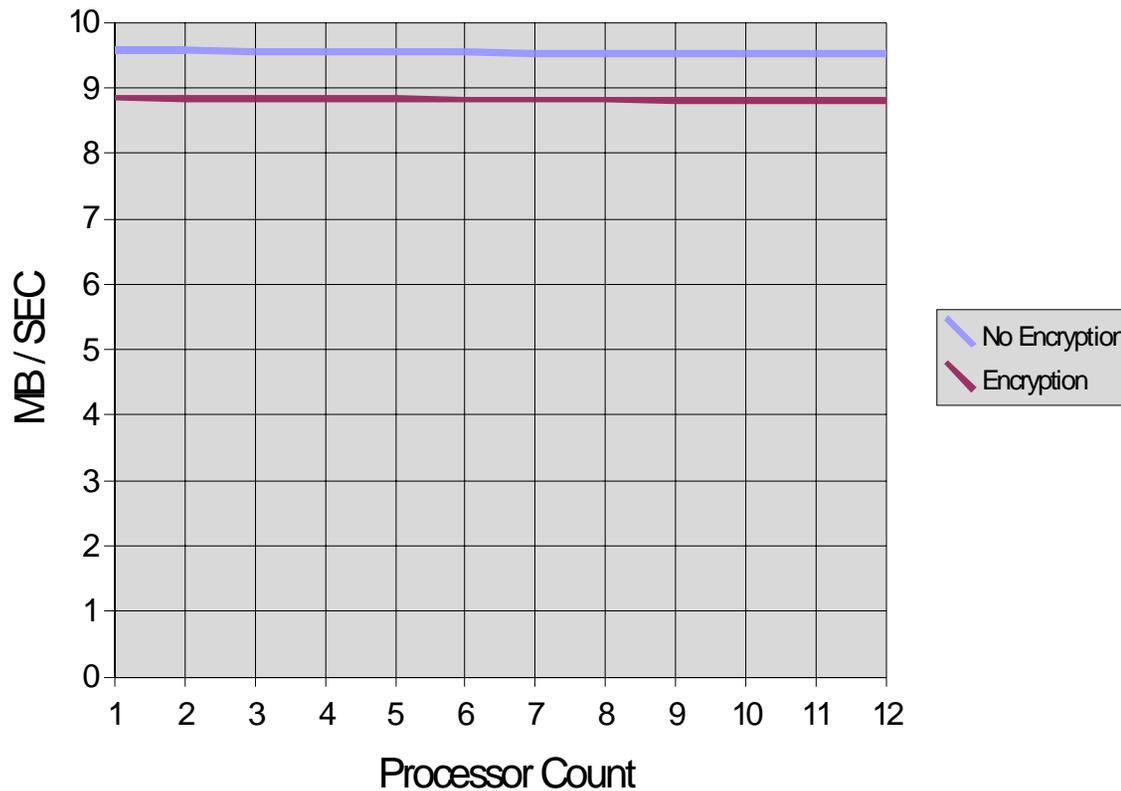


Transactions Per Second



Openais – performance with many processors

Group Messaging Throughput



openais – project statistics

- Executive LOC: 29728
- Library LOC: 9505
- Include LOC: 4231
- Total LOC: 43464
- Changesets since openais inception: 880
- estimated 4500 hours of development, 9000 hours of testing and QA

Accomplishments 2005

Distributed locks API merged

Integration with Linux-HA

Integration with Greenplum frame support, 70MB/sec throughput

Initial Replicant Ring protocol

Patrick Caulfield submitted Ipv6 support

AMF B.01.01 initial patch

98% pass rate with SAFtest

3rd Generation Protocol implemented (totem)

Jan

Apr

June

July

Au

Sister Projects

- The totem protocol in openais integrated into Redhat's CMAN
- The EVS library used by Linux-HA to support membership and messaging
- Integration of openais's AMF, CKPT with Asterisk as POC
- openais has generated 85 bug fixes (patches) for safest AIS B

roadmap

Picacho
AMF
service
CKPT

service

EVT

service

CLM

service

EVS

B.01.01

B.01.01

85% code
coverage

SAF test

95% pass
rate targeted
for

december

INCREASE 1

plus:

- DLCK
service
(exp)

- MSG
service
(exp)

- AMF
B.01.01

- Prototype

redundant
ring

- AMF

management
Q1/06

- Targeted

S

Release 1
plus:

- DLCK
service

- MSG
service

- AMF

B.01.01

- Production

redundant
ring

- AMF

management
Q4/05

- Targeted

Conclusion

- Reduce MTTR to improve availability
- SA Forum AIS provides APIs to reduce MTTR
- open source solution available of AIS (<http://developer.osdl.org/dev/openais>)
- openais is suitable for deployment today

Questions?