

The openais project

Technomap

Prepared by
Steven Dake
January 2005

Agenda

- 2004 Accomplishments
- Current & Future Technology
- Technomap
- CGL Cluster Spec Analysis

2004 Accomplishments

3rd generation protocol
Implemented and merged

First project users

OSDL and SA Forum
Announce official support

EVS service merged

AMF Configuratioin
Support merged by
Sakai Miyotaka

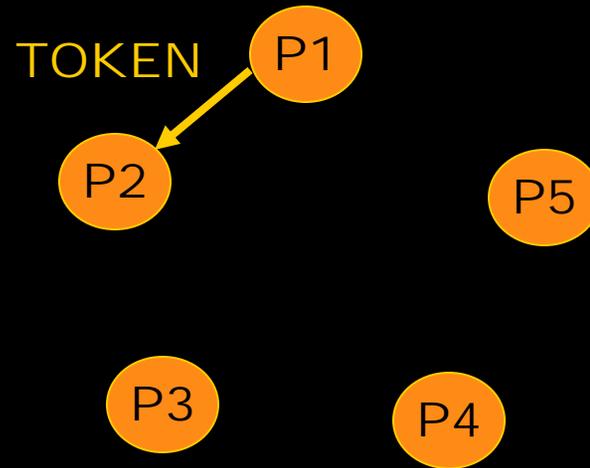
EVT service merged
By Mark Haverkamp

Pthreads library support

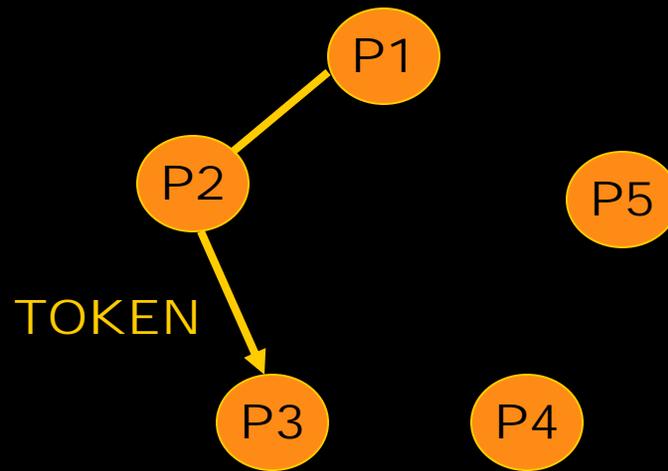
Project open sourced
Included AMF, CLM, CKPT
Using virtual synchrony

July August September October November December January

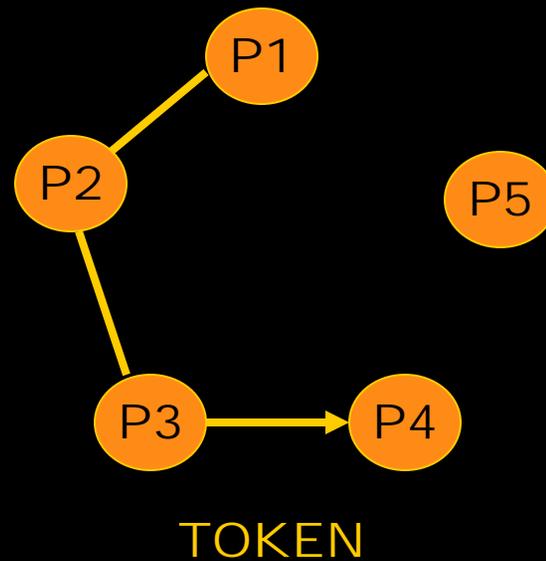
Current Technology - Single Ring Protocol



Current Technology - Single Ring Protocol



Current Technology - Single Ring Protocol



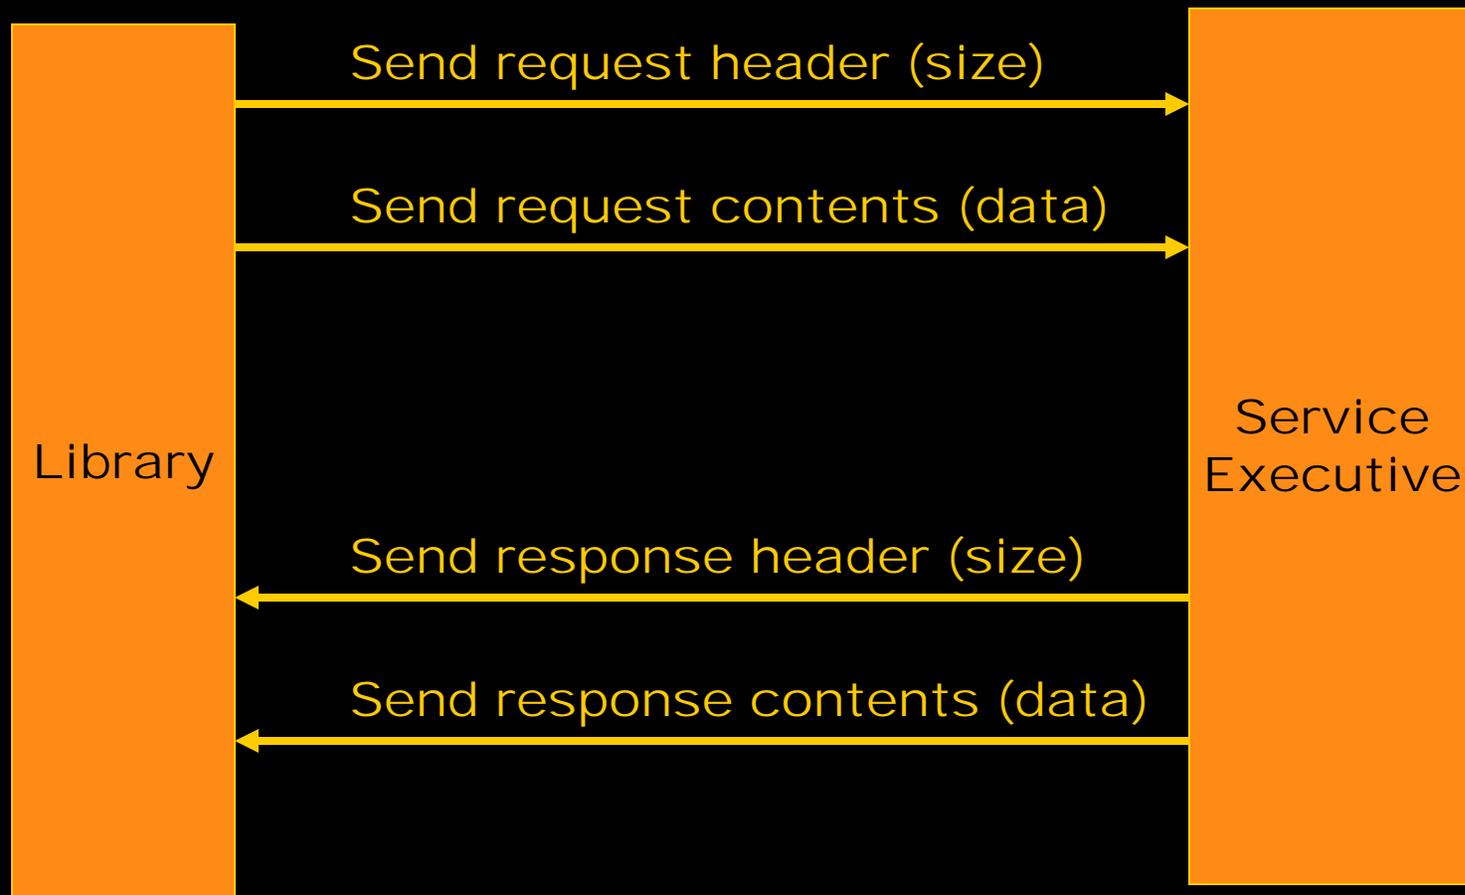
Single Ring Protocol – Problems

- Latency from origination to agreed deliver is $\frac{1}{2}$ token rotation time
- Bandwidth per processor is total/processor count when all processors transmit maximum

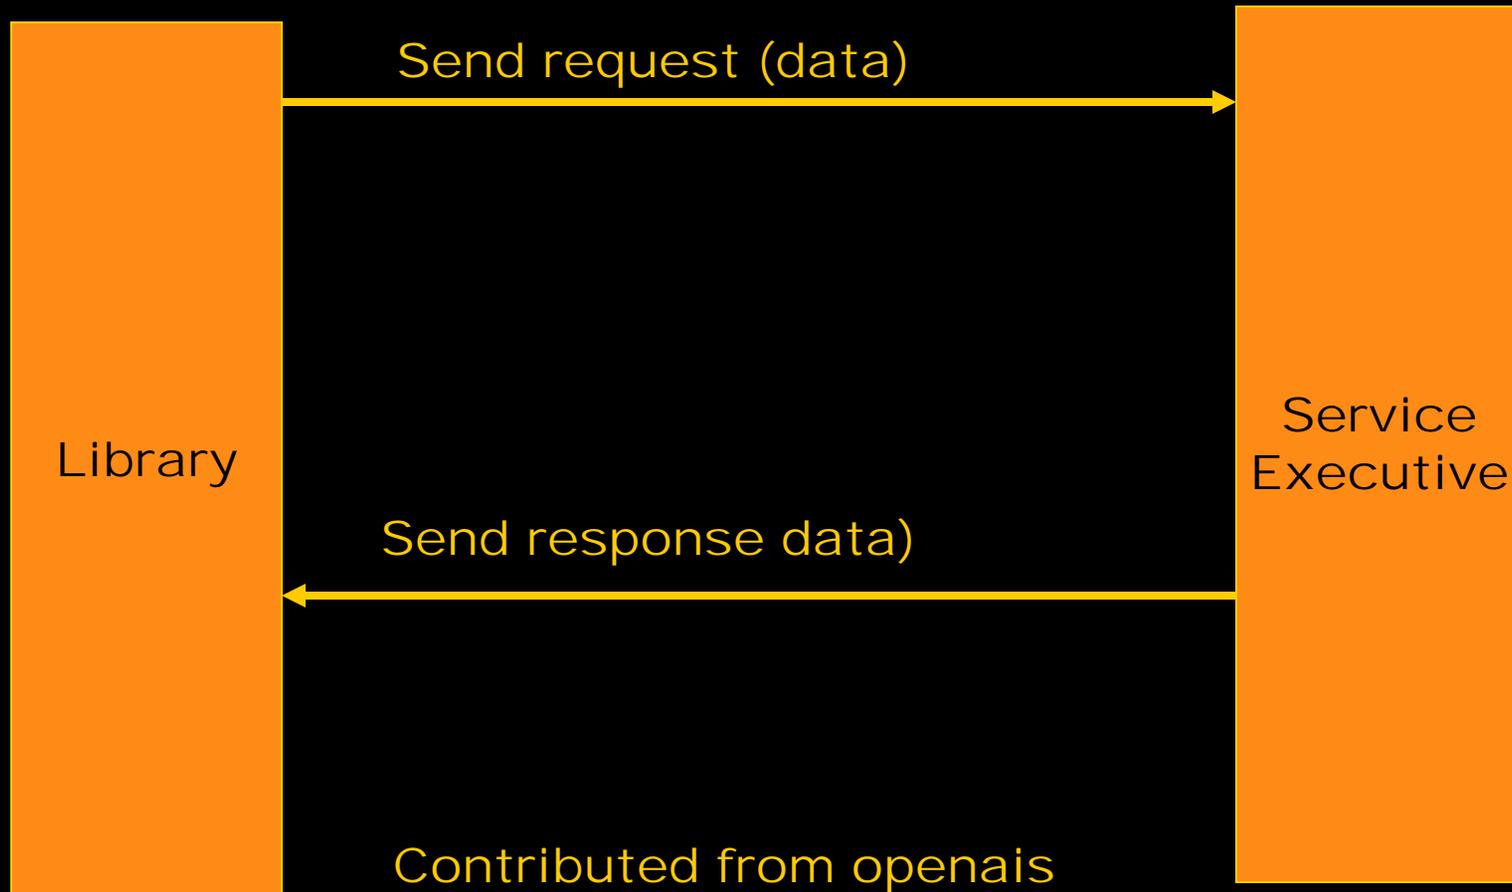
Future Technology - Multi Ring Protocol



Current Technology - IPC System



Future Technology - SOCK_SEQPACKET



Contributed from openais
Project to linux 2.6.3

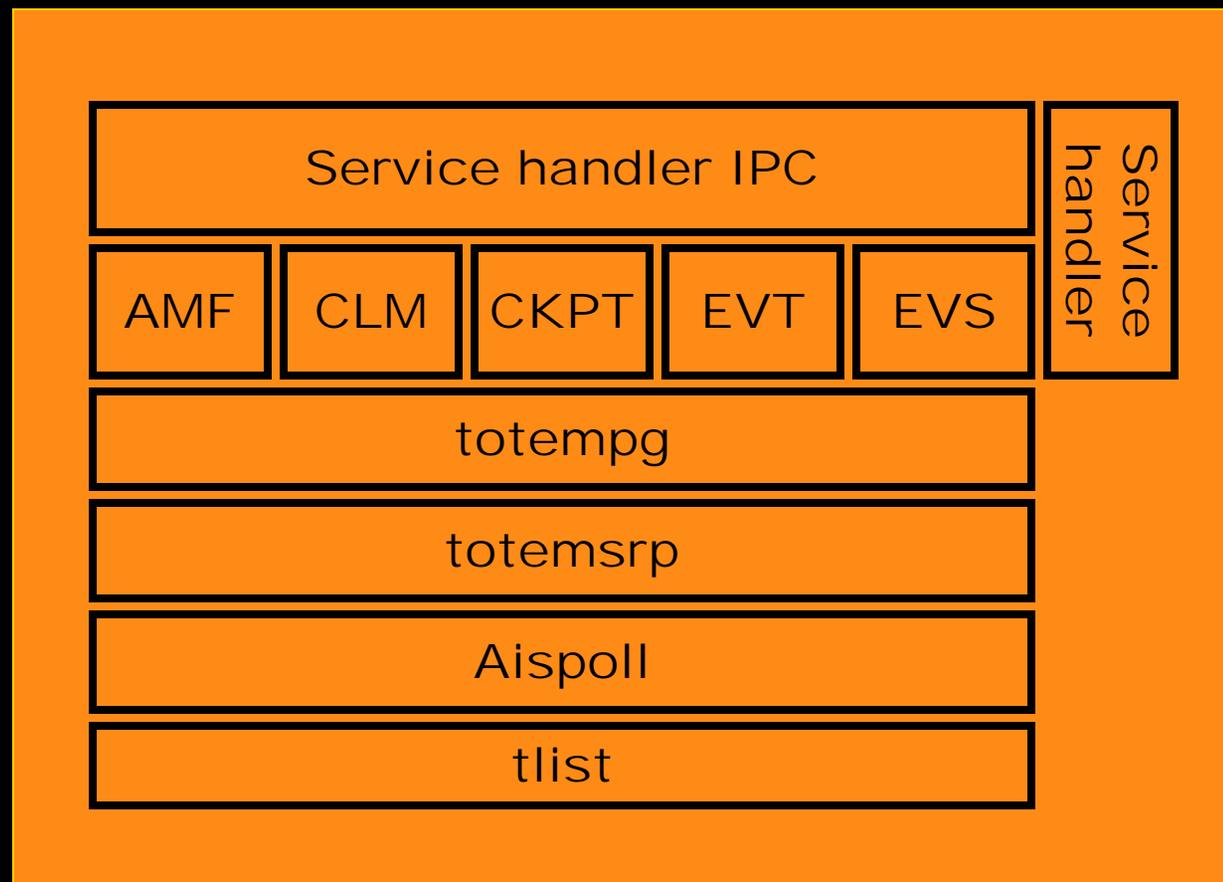
Current Technology – Strings

- Strings stored in 256 byte arrays
- String sent in almost every request within AIS
- Small requests = lots of overhead

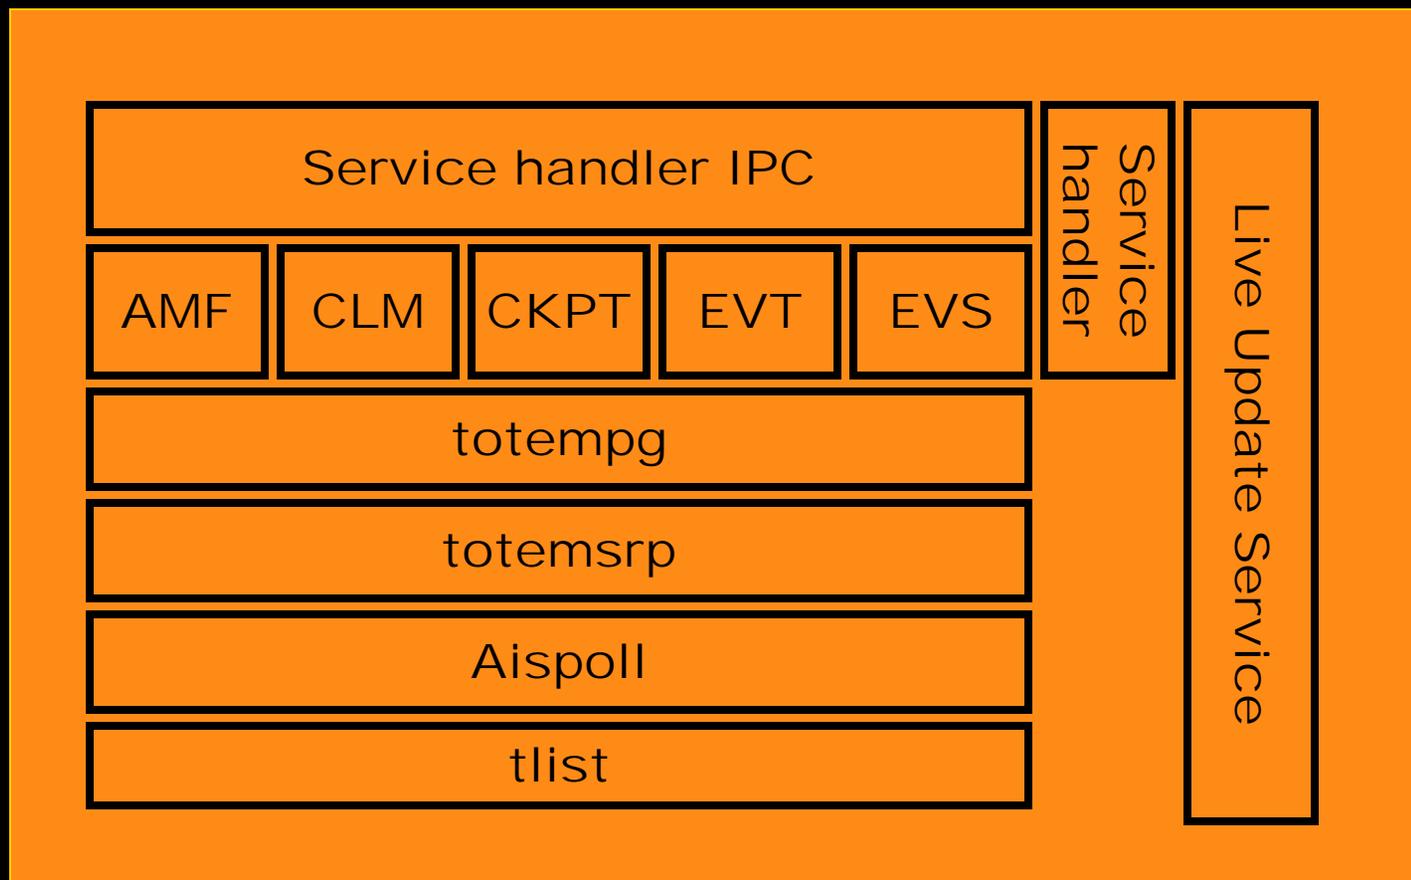
New Technology - Distributed Name Service

- 32 bit handle used to reference a name
- Name/string database synchronized after partition change
- Lookup of handle is $O(1)$
- Lookup of string uses hash table

Current Architecture



Future Technology - Live Update



Current Technology - Secrecy & Authentication

- HMAC/SHA1 used to authenticate messages
- SOBER128 used to encrypt data
- Current weakness is private key secrecy, lack of rekeying

New Technology – Group Key Generation

- Choose one of the 6 published group key algorithms and implement
- Private key stored on each individual processor
- Public key for all processors stored on all processors
- Key used for encryption and authentication generated by math applied to public and private keys within group
- Supports arbitrary rekeying

New Technology - Manageability

- Complete lack of external management of AMF components
- Perhaps a web interface to manage components and provide full AMF semantics

And of course priority #1

- Implement message service
- Implement distributed lock service
- Transition AMF to B.01.01

The openais technomap

Release 1

- AMF service
- CKPT service
- EVT service
- CLM service
- EVS service
- B.01.01
- 85% code coverage
- SAFTEST run

Release 2

Release 1 plus:

- DLCK service
- MSG service
- RMD service
- AMF B.01.01
- Prototype multiring
- SEQPACKET
- AMF management
- Distributed name service

Release 3

Release 2 plus:

- Update to current spec
- Production multiring
- Live update
- Group key generation
- 90% code coverage

Q1/05

Q4/05

Q4/06

CGL Analysis - We will do these:

- CMS 1.0 - Cluster Membership Service
- CES.1.0 - Cluster Event Service
- CCS.1.0 Cluster Checkpoint Service
- CCM.1.0 Cluster Message Service
- CLS 1.0 Cluster Lock Service
- CAF1.0 Cluster Availability Framework
- CFH.1.0 Cluster Node Detection
- CFH.3.0 Application Failover Enabling
- CCM.2 Cluster Communication Service
- CAF.2.1 Ethernet MAC Address Takeover
- CAF.2.2 IP Takeover
- CCS.2.0 SAF_AIS Data Checkpoint Performance
- CMS.2.0 Dynamic Cluster Membership
- CCON.1.1 Run Diagnostics
- CCON.1.2 Boot Reboot Nodes
- CCM.4.1/CCM4.3 Group Messaging Protocol

CGL Analysis - We wont do these (anytime soon):

- CFH.2.0 Prevent Failed Node from Corrupting Shared Resources
- CSM.1.0 Storage Network Replication
- CSM.2.0 Cluster aware Volume Management
- CSM.3.0 Shared Storage Mirroring
- CSM.4.0 Redundant Cluster Storage Path
- CSM.5.0 Cluster Synchronized Device Hotswap
- CSM.6.0 Cluster Filesystem
- CAF.2.3 Deliberate TCP Session Takeover
- CMON.1 Cluster Node Monitoring
- CCON.1.1 Run Diagnostics
- CDIAG.1.1/1.2 Online Diagnostics
- CDIAG.2.1/2.4 Cluster Wide Diagnostic Info

Conclusions

- The openais project is making good progress
- The project has a competitive roadmap
- We are not the cure to missing implementations of requirements in CGL

Join in

- WEB
<http://developer.osdl.org/dev/openais>
- Mailing List
openais@lists.osdl.org

We need docs, tests, code, designs.

There is still A LOT of design work to do come join in on the early stages of the project.