

Common API for Transparent Hybrid Multicast draft-irtf-samrg-common-api - Status Update –

Matthias Wählisch, Thomas C. Schmidt
Stig Venaas

{waelisch, t.schmidt}@ieee.org,
stig@cisco.com

History of the Draft

- o Version 00/01 presented @ IETF 76, Hiroshima (Nov. 2009)
- o Adopted as WG document @ IETF 79, Beijing (Nov. 2010)
- o Update version 03: Major update on namespace and mapping + Many clarifications and editorial improvements
- o Currently published version: 04
 - Considers comments by WeeSan and Mario Kolberg
 - Section “Illustrative Example” + “Implementation” added
- o RGLC February 2012
 - Several comments by John and Sebastian
 - Discussions about MTU issues

Comments by John (1)

- o Please add language that specifies whether group names are intended to be unique (by namespace domain, by all domains, across all active sessions, across all past present or future sessions, other)?
 - Group Name is globally unique by all domains
- o Is the application responsible for selecting group names and determining uniqueness?
 - Group Names can be chosen by the end user, which component ensures uniqueness is out of scope of the document
- o Clarification of service discovery: Implementation-specific, addressed in another spec., or use existing mechanisms

Comments by John (2)

- o Scheme extension: Is “sips” a possible scheme?
 - Yes.
- o Is the `receive` call blocking?
 - Yes, in the reference implementation.
- o How is mobility across multicast domains supported?
 - Hiding mobility by Group Names
- o Support of multi-resolution multicast and XCAST
 - Next slides ...

How does the API cover multi-resolution Multicast?

- o General objective of multi-resolution multicast:
Consider heterogeneous end devices
- o Idea: Utilize the mapping from names to addresses
- o Option 1: Different *addresses* for different quality
 - Domain-specific mapping
- o Option 2: Pre-defined *naming* syntax
 - Example for a layered video:
`ELj.Qi.blockbuster` denotes each individual layer

Support of XCAST?

- o XCAST idea: Source explicitly pre-defines receiver
- o Perspective of the API
 - XCAST is another distribution technology: new Interface
 - scheme://group@instantiation – instantiation denotes one or multiple *sources* – one may encode receivers using an explicit XCAST scheme
 - Receiver Ids are technology specific which would contradict abstraction approach of names

Comments by Sebastian

- o At the moment error handling is quite limited and just indicates success or failure of send/recv operation. From a developers perspective this might be a little coarse, as there is no information on what to do about the error. A common problem would be the **size of data (message)**, that an application would like to send within one send call.
 - Next version of the draft includes error codes for send/receive calls that indicates that message is too long

MTU Issues

- o Current version does not provide mechanisms to handle MTU problems

Challenge:

- o API considers a multi-technology scenario
 - Different technologies have different MTU sizes
 - MTU size may change over time
 - Interfaces may change dynamically (mobility)
- o No more complexity to the application programmer

A **too** Simple Approach

- o Each interface signals currently supported MTU
- o Programmer configures MaxMessageSize with socket creation

Problems:

- o No appropriate interface(s) available
- o Change of path MTU during transmission
 - Interfaces will be disabled
 - Prevents data delivery

Proposal

- o Middleware implementation guarantees a `maxMessageSize` to the application programmer
- o Fragmentation will be implemented by technology modules if necessary
- o No incompatibility to native IP clients due to realistic values for `maxMessageSize`

Consequences - Draft Changes

Call extensions:

- o getMaxMsgSize() as Service Call
- o getAtomicMsgSize() as Socket Option

Editorial extensions:

- o Separate MTU Section that discusses the general problem scope