

Document Title	Explanation of ara::com API
Document Owner	AUTOSAR
Document Responsibility	AUTOSAR
Document Identification No	846

Document Status	published
Part of AUTOSAR Standard	Adaptive Platform
Part of Standard Release	R24-11

Document Change History			
Date	Release	Changed by	Description
2024-11-27	R24-11	AUTOSAR Release Management	<ul style="list-style-type: none"> Updated tutorial examples to use complete ARXML
2023-11-23	R23-11	AUTOSAR Release Management	<ul style="list-style-type: none"> CommunicationGroups is now OBSOLETE Removed Raw Data Stream chapters (moved to AP EXP InterfaceGuidelines)
2022-11-24	R22-11	AUTOSAR Release Management	<ul style="list-style-type: none"> Fully reworked version of this document. Added new Tutorials chapter.
2021-11-25	R21-11	AUTOSAR Release Management	<ul style="list-style-type: none"> No changes. Fully reworked version of this document is going to be released in R22-11.
2020-11-30	R20-11	AUTOSAR Release Management	<ul style="list-style-type: none"> Replaced term "(Un)Checked Exception" by proper formulations Clarified the usage and transference of "Instance Specifier" Removed the reference to "AUTOSAR_RS_CPP14Guidelines"
2019-11-28	R19-11	AUTOSAR Release Management	<ul style="list-style-type: none"> Added access to current field value from Get/SetHandler Changed Document Status from Final to published
2019-03-29	19-03	AUTOSAR Release Management	<ul style="list-style-type: none"> Changed explanation of Event reception due to new ara::com API



△

2018-10-31	18-10	AUTOSAR Release Management	<ul style="list-style-type: none"> ● Added InstanceIdentifier and InstanceSpecifier explanation ● Restructured chapter structure ● Adapt FindService signatures ● Added sample code for event usage ● Restructured chapter structure ● Proxy and skeleton instances are not copyable ● Changed certain data types to ara::core namespace. ● Adapted to new error handling based on ara::core::ErrorCode
2018-03-29	18-03	AUTOSAR Release Management	<ul style="list-style-type: none"> ● Added Fire&Forget Methods ● Minor changes and bugfixes
2017-10-27	17-10	AUTOSAR Release Management	<ul style="list-style-type: none"> ● Added explanation of TLV ● Minor changes and bugfixes
2017-03-31	17-03	AUTOSAR Release Management	<ul style="list-style-type: none"> ● Initial release

Disclaimer

This work (specification and/or software implementation) and the material contained in it, as released by AUTOSAR, is for the purpose of information only. AUTOSAR and the companies that have contributed to it shall not be liable for any use of the work.

The material contained in this work is protected by copyright and other types of intellectual property rights. The commercial exploitation of the material contained in this work requires a license to such intellectual property rights.

This work may be utilized or reproduced without any modification, in any form or by any means, for informational purposes only. For any other purpose, no part of the work may be utilized or reproduced, in any form or by any means, without permission in writing from the publisher.

The work has been developed for automotive applications only. It has neither been developed, nor tested for non-automotive applications.

The word AUTOSAR and the AUTOSAR logo are registered trademarks.

Contents

1	Preface	9
2	Acronyms and Abbreviations	10
3	Introduction	13
3.1	Approach	13
3.2	API Design Visions and Guidelines	15
4	Fundamentals	17
4.1	Proxy/Skeleton Architecture	17
4.2	Means of Communication	18
4.3	ara::com Event and Trigger based communication	18
4.4	ara::com Method based communication	19
4.5	ara::com Field based communication	20
4.6	Data Type Abstractions	21
4.7	Error Handling	22
4.8	Service Connection Approach	22
4.8.1	Instance Identifiers and Instance Specifiers	22
4.8.2	When to use InstanceIdentifier versus InstanceSpecifier	25
4.8.2.1	Transfer of an InstanceIdentifier	26
5	Detailed API description	27
5.1	High Level API Structure	27
5.2	API Elements	27
5.3	Proxy Class	29
5.3.1	Proxy Class API's	30
5.3.2	RadarService Proxy Class Example	30
5.3.3	Constructor and Handle Concept	33
5.3.4	Finding Services	34
5.3.4.1	Auto Update Proxy instance	35
5.3.5	Events	38
5.3.5.1	Event Subscription and Local Cache	41
5.3.5.2	Monitoring Event Subscription	42
5.3.5.3	Accessing Event Data — aka Samples	45
5.3.5.4	Event Sample Management via <code>SamplePtrs</code>	46
5.3.5.5	Event-Driven vs Polling-Based access	47
5.3.5.6	Buffering Strategies	49
5.3.6	Methods	52
5.3.6.1	One-Way aka Fire-and-Forget Methods	54
5.3.6.2	Event-Driven vs Polling access to method results	55
5.3.6.3	Canceling Method Result	59
5.3.7	Fields	60
5.3.8	Triggers	63
5.4	Skeleton Class	63

5.4.1	Skeleton Class API's	64
5.4.2	RadarService Skeleton Class Example	64
5.4.3	Instantiation (Constructors)	67
5.4.4	Offering Service instance	68
5.4.5	Polling and event-driven processing modes	69
5.4.5.1	Polling Mode	70
5.4.5.2	Event-Driven Mode	71
5.4.6	Methods	73
5.4.6.1	One-Way aka Fire-and-Forget Methods	76
5.4.6.2	Raising Application Errors	77
5.4.7	Events	78
5.4.8	Fields	81
5.4.8.1	Registering Getters	82
5.4.8.2	Registering Setters	83
5.4.8.3	Ensuring existence of "SetHandler"	84
5.4.8.4	Ensuring existence of valid Field values	84
5.4.8.5	Access to current field value from Get/SetHandler	84
5.4.9	Triggers	84
5.5	Data Types on Service Interface level	85
5.5.1	Optional data elements	85
5.6	Communication Groups	88
5.6.1	Objective	88
5.6.2	Realization	89
5.6.3	Usage Scope	90
5.6.4	Special Topics	91
5.6.4.1	Architecture	91
5.6.4.2	Remote Connection	92
5.6.4.3	Service Versioning	92
6	Tutorials	93
6.1	Usage of Service Interfaces	93
6.1.1	Service Interface Deployment	93
6.1.2	Service Instance Deployment	94
6.1.3	Service Implementation	98
6.2	Usage of InstanceSpecifier	99
6.2.1	Modeling and configuration/mapping over Manifest from user perspective	101
6.2.2	Instance IDs only for provided Services	106
6.3	Usage in context of MultiBinding	106
6.4	Usage of CommunicationGroups	113
6.4.1	Setup	114
6.4.1.1	Service Interface Description Generation	114
6.4.1.2	Proxy and Skeleton Generation	115
6.4.2	Example	117
7	Appendix	121
7.1	Serialization	121

7.1.1	Zero-Copy implications	122
7.2	Service Discovery Implementation Strategies	122
7.2.1	Central vs Distributed approach	123
7.3	Multi-Binding implications	126
7.3.1	Simple Multi-Binding use case	126
7.3.2	Local/Network Multi-Binding use case	129
7.3.3	Typical SOME/IP Multi-Binding use case	130
7.4	ara::com and AUTOSAR meta-model relationship	132
7.4.1	Connection to AUTOSAR_TR_AdaptiveMethodology	133
7.4.2	Service Interface	133
7.4.3	Software Component	134
7.4.4	Adaptive Application/Executables and Processes	136
7.4.5	Usage of meta-model identifiers within ara::com based application code	137
7.5	Abstract Protocol Network Binding Examples	139

References

- [1] Specification of Manifest
AUTOSAR_AP_TPS_ManifestSpecification
- [2] Specification of Adaptive Platform Core
AUTOSAR_AP_SWS_Core
- [3] Specification of Communication Management
AUTOSAR_AP_SWS_CommunicationManagement
- [4] Specification of State Management
AUTOSAR_AP_SWS_StateManagement
- [5] Software Component Template
AUTOSAR_CP_TPS_SoftwareComponentTemplate
- [6] Specification of RTE Software
AUTOSAR_CP_SWS_RTE
- [7] Middleware for Real-time and Embedded Systems
<http://doi.acm.org/10.1145/508448.508472>
- [8] Patterns, Frameworks, and Middleware: Their Synergistic Relationships
<http://dl.acm.org/citation.cfm?id=776816.776917>
- [9] SOME/IP Protocol Specification
AUTOSAR_FO_PRS_SOMEIPProtocol
- [10] E2E Protocol Specification
AUTOSAR_FO_PRS_E2EProtocol
- [11] Serialization and Unserialization
<https://isocpp.org/wiki/faq/serialization>
- [12] Copying and Comparing: Problems and Solutions
http://dx.doi.org/10.1007/3-540-45102-1_11
- [13] SOME/IP Service Discovery Protocol Specification
AUTOSAR_FO_PRS_SOMEIPServiceDiscoveryProtocol

1 Preface

Typically, reading formal specifications isn't the easiest way to learn and understand a certain technology. This especially holds true for the Communication Management API ([ara::com](#)) in the AUTOSAR Adaptive Platform.

Therefore this document shall serve as an entry point not only for the developer of software components for the Adaptive Platform, who will use the [ara::com](#) API to interact with other application or service components, but also for Adaptive Platform product vendors, who are going to implement an optimized `IPC` binding for the [ara::com](#) API on their platform.

We strongly encourage both groups of readers to read this document at hand before going into the formal details of the related SWS.

2 Acronyms and Abbreviations

The glossary below includes acronyms and abbreviations relevant to the explanation of ara::com API.

Abbreviation / Acronym:	Description:
ara::com	C++ namespace of functional cluster Communication Management
COMMUNICATION_GROUP_SERVER	COMMUNICATION_GROUP_SERVER according to [TPS-MANI_03628] of [1]
COMMUNICATION_GROUP_CLIENT	COMMUNICATION_GROUP_CLIENT according to [TPS-MANI_03628] of [1]
ctor	C++ constructor
dtor	C++ destructor
std::future	C++ std::future provides a mechanism to access the result of asynchronous operations
std::promise	C++ std::promise provides a facility to store a value or an exception that is later acquired asynchronously via a std::future object created by the std::promise object
Error	Recoverable Errors according to [SWS_CORE_00020] of [2]. Defined for ara::com functions using ara::core::ErrorCodes, e.g. according to [SWS_CM_10432] of [3]
ara::core::ErrorCode	ara::core::ErrorCode according to [SWS_CORE_00501] of [2]
Violation	Violation according to [SWS_CORE_00021] of [2]
Corruption	Corruption according to [SWS_CORE_00022] of [2]
ara::com::InstanceIdentifier	ara::com Instance Identifier according to [SWS_CM_00302] of [3]
ara::com::InstanceIdentifierContainer	ara::com Instance Identifier Container according to [SWS_CM_00319] of [3]
ara::com::ComErrorDomain	ara::com::ComErrorDomain according to [SWS_CM_10432] of [3]
ara::com::e2e::E2EErrorDomain	ara::com::e2e::E2EErrorDomain according to [SWS_CM_10474] of [3]
ara::core::Result	Returned result object according to [SWS_CORE_00701] of [2]
ara::core::Future	Returned future object according to [SWS_CORE_00322] of [2]
ara::core::Promise	Returned promise object according to [SWS_CORE_00341] of [2]
ara::core::Optional	Provides access to optional record elements according to [SWS_CORE_01033] of [2]
ara::core::GetResult	GetResult according to [SWS_CORE_00336] of [2]
ara::com SubscriptionState	SubscriptionState according to [SWS_CM_00310] of [3]
ara::com MethodCallProcessingMode	MethodCallProcessingMode according to [SWS_CM_00301] of [3]
ara::core::InstanceSpecifier	ara::core Instance Specifier according to [SWS_CORE_08001] and following of [2]
ara::com::HandleType	ara::com HandleType according to [SWS_CM_00312] and following of [3]
ara::com::FindServiceHandle	ara::com FindServiceHandle according to [SWS_CM_00303] and following of [3]
ara::com::FindServiceHandler	ara::com FindServiceHandler according to [SWS_CM_00383] and following of [3]
ara::com::SamplePtr	ara::com::SamplePtr according to [SWS_CM_00306] and following of [3]
ara::com::SampleAllocateePtr	ara::com::SampleAllocateePtr according to [SWS_CM_00308] and [SWS_CM_00306] and following of [3]

Abbreviation / Acronym:	Description:
Proxy::Event::Subscribe	EventSubscribe according to [SWS_CM_00141] of [3]
Proxy::Trigger::Subscribe	TriggerEventSubscribe according to [SWS_CM_00723] of [3]
Proxy::Event::Unsubscribe	EventUnsubscribe according to [SWS_CM_00151] of [3]
Proxy::Trigger::Unsubscribe	TriggerUnsubscribe according to [SWS_CM_00151] of [3]
Proxy::GetSubscriptionState	GetSubscriptionState according to [SWS_CM_00316] of [3]
Proxy::SetSubscriptionStateHandler	GetSubscriptionState according to [SWS_CM_00333] of [3]
Proxy::SetSubscriptionStateHandler with ExecutorT	SetSubscriptionState with ExecutorT according to [SWS_CM_11354] of [3]
Proxy::UnsetSubscriptionStateChangeHandler	UnsetSubscriptionStateChangeHandler according to [SWS_CM_00334] of [3]
Proxy::GetNewSamples	GetNewSamples according to [SWS_CM_00701] of [3]
Proxy::GetFreeSampleCount	GetFreeSampleCount according to [SWS_CM_00705] of [3]
Proxy::SetReceiveHandler	SetReceiveHandler according to [SWS_CM_00181] of [3]
Proxy::SetReceiveHandler with ExecutorT	SetReceiveHandler with ExecutorT according to [SWS_CM_11356] of [3]
Proxy::UnsetReceiveHandler	UnsetReceiveHandler according to [SWS_CM_00183] of [3]
Proxy::ResolveInstanceIDs	ResolveInstanceIDs according to [SWS_CM_00118] of [3]
Proxy::Field Get	FieldGet according to [SWS_CM_00112] of [3]
Proxy::Field Set	FieldSet according to [SWS_CM_00113] of [3]
Proxy::FindService	FindService according to [SWS_CM_00622] of [3]
Proxy::FindService with Instance Specifier	FindService with Instance Specifier according to [SWS_CM_00623] of [3]
Proxy::StartFindService	StartFindService according to [SWS_CM_00123] of [3]
Proxy::StartFindService with ExecutorT	StartFindService with ExecutorT according to [SWS_CM_11352] of [3]
Proxy::StopFindService	StopFindService according to [SWS_CM_00125] of [3]
Skeleton::OfferService	OfferService according to [SWS_CM_00101] of [3]
Skeleton::StopOfferService	StopOfferService according to [SWS_CM_00111] of [3]
Skeleton::Trigger Send	Trigger::Send according to [SWS_CM_00721] of [3]
Skeleton::Event::Send	Event::Send according to [SWS_CM_00162] of [3]
Skeleton::Allocate	Allocate according to [SWS_CM_90438] of [3]
Skeleton::Send with SampleAllocateePtr	Send with SampleAllocateePtr according to [SWS_CM_90437] of [3]
Skeleton::ProcessNextMethodCall	ProcessNextMethodCall according to [SWS_CM_00199] of [3]
Skeleton::RegisterGetHandler	RegisterGetHandler according to [SWS_CM_00114] of [3]
Skeleton::RegisterGetHandler with ExecutorT	RegisterGetHandler with ExecutorT according to [SWS_CM_11360] of [3]
Skeleton::RegisterSetHandler	RegisterSetHandler according to [SWS_CM_00116] of [3]
Skeleton::RegisterSetHandler with Executort	RegisterSetHandler with Executor according to [SWS_CM_11362] of [3]
Skeleton::Field Update	Update according to [SWS_CM_00119] of [3]
IPC	Inter Process Communication
RT	Realtime
SI	Service Interface
WET	Worst Case Execution Time
PowerMode	PowerMode according to [SWS_SM_91020] of [4]
FieldSenderComSpec	FieldSenderComSpec according to [TPS_MANI_03211] of [1]
PPortPrototype	PPortPrototype according to [TPS_SWCT_01111] of [5]
CppImplementationDataType	CppImplementationDataType according to [TPS_MANI_01166] of [1]

Terms:	Description:
Binding	<p>This typically describes the realization of some abstract concept with a specific implementation or technology.</p> <p>In AUTOSAR, for instance, we have an abstract data type and interface model described in the methodology.</p> <p>Mapping it to a concrete programming language is called <i>language binding</i>. In the AUTOSAR Adaptive Platform for instance we do have a C++ language binding.</p> <p>In this explanatory document we typically use the tech term <i>binding</i> to refer to the implementation of the abstract (technology independent) ara::com API to a concrete communication transport technology like for instance sockets, pipes, shared memory, ...</p>
Callable	<p>In the context of C++ a Callable is defined as: A Callable type is a type for which the INVOKE operation (used by, e.g., std::function, std::bind, and std::thread::thread) is applicable. This operation may be performed explicitly using the library function std::invoke. (since C++17)</p>

3 Introduction

3.1 Approach

Why did AUTOSAR invent yet another communication middleware API/technology, while there are dozens on the market — the more so as one of the guidelines of Adaptive Platform was to reuse existing and field proven technology?

Before coming up with a new middleware design, we did evaluate existing technologies, which — at first glance — seemed to be valid candidates. Among those were:

- ROS API
- DDS API
- CommonAPI (GENIVI)
- DADDY API (Bosch)

The final decision to come up with a new and AUTOSAR-specific Communication Management API was made due to the fact, that not all of our key requirements were met by existing solutions:

- We need a Communication Management, which is NOT bound to a concrete network communication protocol. It has to support the SOME/IP protocol but there has to be flexibility to exchange that.
- The AUTOSAR service model, which defines services as a collection of provided methods, events and fields shall be supported naturally/straight forward.
- The API shall support an event-driven and a polling model to get access to communicated data equally well. The latter one is typically needed by real-time applications to avoid unnecessary context switches, while the former one is much more convenient for applications without real-time requirements.
- Possibility for seamless integration of end-to-end protection to fulfill ASIL requirements.
- Support for static (preconfigured) and dynamic (runtime) selection of service instances to communicate with.

So in the final `ara::com` API specification, the reader will find concepts (which we will describe in-depth in the upcoming chapters), which might be familiar for him from technologies, we have evaluated or even from the existing Classic Platform:

- Proxy (or Stub)/Skeleton approach (CORBA, Ice, CommonAPI, Java RMI, ...)
- Protocol independent API (CommonAPI, Java RMI)
- Queued communication with configurable receiver-side caches (DDS, DADDY, Classic Platform)

- Zero-copy capable API with possibility to shift memory management to the middleware (DADDY)
- Data reception filtering (DDS, DADDY)

Now that we have established the introduction of a new middleware API, we go into the details of the API in the following chapters.

The following statement is the basis for basically all AUTOSAR AP specifications, but should be explicitly pointed out here again:

ara::com only defines the API signatures and its behavior visible to the application developer. Providing an implementation of those APIs and the underlying middleware transport layer is the responsibility of the AUTOSAR AP vendor.

For a rough parallel with the AUTOSAR Classic Platform, `ara::com` can be seen as fulfilling functional requirements in the Adaptive Platform similar to those covered in the Classic Platform by the RTE APIs [6] such as `Rte_Write`, `Rte_Read`, `Rte_Send`, `Rte_Receive`, `Rte_Call`, `Rte_Result`.

Overview of Modeling elements and how they are related to each other: [SI](#), Deployment, Actual generation dependant from provided Deployment Information (E.g. also [SI](#) Elements that will be generated later and connection to Service Instance Manifest)

AUTOSAR Adaptive Platform methodology explains the process aspects necessary to build an Adaptive AUTOSAR system and how they relate to each other [TR_AMETH_00100]. It defines activities and work products delivered or consumed [TR_AMETH_00102] and the Roles performed by OEMs and suppliers.

Major steps involved in the development of Adaptive Software are

- Architecture and Design
- Adaptive Software Development
- Integration and Deployment

Adaptive applications run on top of ARA layer and exchanges the information using [SIs](#) and Ports. Important contribution for `ara::com` API work performed during the Integration and Deployment step of Adaptive Methodology. It supports the generation of [SI](#) Description ARXML file, which aggregates the [SIs](#) and ports. [SIs](#) for service-oriented communication defined by Events, Methods and Fields [Listing 5.1]. This is done independent of Software components or Transport layer used for underlying communication.

Adaptive Platform supports two types of ports namely Provided and Required. [SI](#) along with Provided port details used for the generation of the Service Skeleton class and Required port details used for the generation of Proxy classes [Figure 5.2]. Proxy and Skeleton classes use `ara::com` API to communicate with other Adaptive Platform clusters and Adaptive Applications.

Service instances are configured, notably the binding of the `SIs` to a chosen transport layer, whether a specific service instance is either Provided or Required and whether there is a mapping to a dedicated Machine. The configurations of the service instance are manifested in the Service Instance Manifest.

Executable of an Adaptive Software are instantiated by means of the Execution Manifest. Instantiation here means to bind the executables to the context of specific processes of the operating system. Each process may start with a different start-up configuration depending on a machine mode. Further on, the Execution Manifest also defines Software process dependencies.

3.2 API Design Visions and Guidelines

One goal of the API design was to have it as lean as possible. Meaning, that it should only provide the minimal set of functionality needed to support the service based communication paradigm consisting of the basic mechanisms: methods, events and fields.

Our definition of the notion "as lean as possible" in this context means: Essentially the API shall only deal with the functionality to handle method, field and event communication on service consumer and service provider implementation side.

If we decided to provide a bit more than just that, then the reason generally was *"If solving a certain communication-related problem ABOVE our API could not be done efficiently, we provide the solution as part of ara::com API layer."*

Consequently, `ara::com` does not provide any kind of component model or framework, which would take care of things like component life cycle, management of program flow or simply setting up `ara::com` API objects according to the formal component description of the respective application.

All this could be easily built on top of the basic `ara::com` API and needs not be standardized to support typical collaboration models.

During the design phase of the API we constantly challenged each part of our drafts, whether it would allow for efficient `IPC` implementations from AP vendors, since we were aware, that you could easily break it already on the API abstraction level, making it hard or almost impossible to implement a well performing binding.

One of the central design points was — as already stated in the introduction — to support polling and event-driven programming paradigms equally well.

So you will see in the later chapters, that the application developer, when using `ara::com` is free to chose the approach, which fits best to his application design, independent whether he implements the service consumer or service provider side of a communication relation.

This allows for support of strictly real-time scheduled applications, where the application requires total control of what (amount) is done when and where unnecessary context switches are most critical.

On the other hand the more relaxed event based applications, which simply want to get notified whenever the communication layer has data available for them is also fully supported.

The decision within AUTOSAR to genuinely support C++11/C++14 for AP was a very good fit for the `ara::com` API design.

For enhanced usability, comfort and a breeze of elegance `ara::com` API exploits C++ features like smart pointers, template functions and classes, proven concepts for asynchronous operations and reasonable operator overloading.

4 Fundamentals

4.1 Proxy/Skeleton Architecture

If you've ever had contact with middleware technology from a programmer's perspective, then the approach of a Proxy/Skeleton architecture might be well known to you.

Looking at the number of middleware technologies using the Proxy/Skeleton (sometimes even called Stub/Skeleton) paradigm, it is reasonable to call it the "classic approach".

So with `ara::com` we also decided to use this classical Proxy/Skeleton architectural pattern and also name it accordingly.

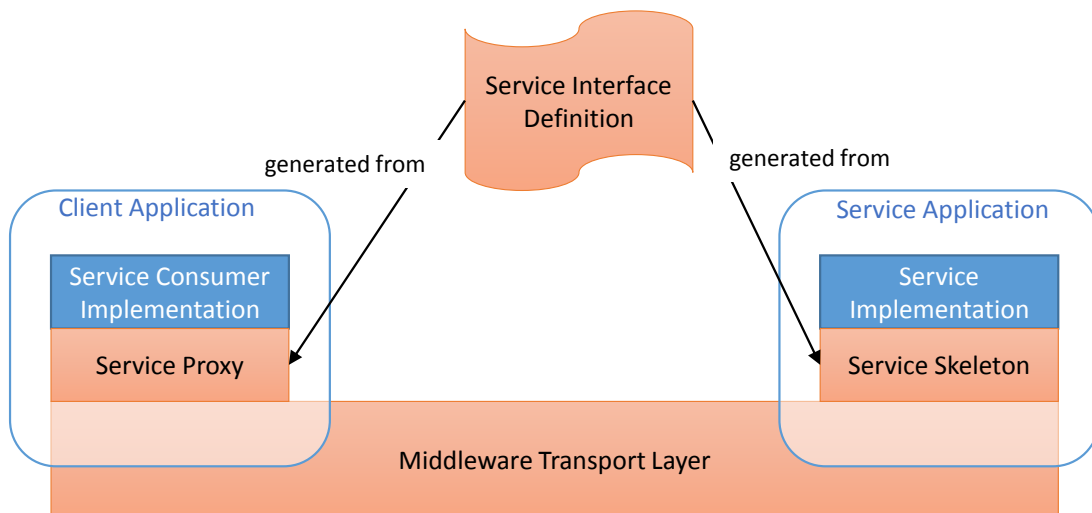


Figure 4.1: Proxy Skeleton Pattern

The basic idea of this pattern is, that from a formal service definition two code artifacts are generated:

- **Service Proxy:** This code is - from the perspective of the service consumer, which wants to use a possibly remote service - the facade that represents this service on code level.

In an object-oriented language binding, this typically is an instance of a generated class, which provides methods for all functionalities the service provides. So the service consumer side application code interacts with this local facade, which then knows how to propagate these calls to the remote service implementation and back.

- **Service Skeleton:** This code is - from the perspective of the service implementation, which provides functionalities according to the service definition - the code, which allows to connect the service implementation to the Communication Man-

agement transport layer, so that the service implementation can be contacted by distributed service consumers.

In an object-oriented language binding, this typically is an instance of a generated class. Usually the service implementation from the application developer is connected with this generated class via a subclass relationship.

So the service side application code interacts with this middleware adapter either by implementing abstract methods of the generated class or by calling methods of that generated class.

Further details regarding the structure of `ara::com` Proxies and Skeletons are shown in section [Section 5.3](#) and [Section 5.4](#). Regarding this design pattern in general and its role in middleware implementations, see [\[7\]](#) and [\[8\]](#).

4.2 Means of Communication

Now, that we've talked about the Proxy/Skeleton Architecture, let us continue to talk about how to communicate between proxies and skeletons.

`ara::com` defines four different mechanisms to communicate between a server and a client

- Methods
- Events
- Fields
- Triggers

Before any of these mechanisms can be used, a service must be instantiated and the server must offer itself to the system (`OfferService()`). Then a client needs to find and connect to the service instance using the Proxy (`FindService()` or `StartFindService()`).

4.3 ara::com Event and Trigger based communication

When a client application has connected to a server, it can subscribe (`Subscribe()`) to events in the service that is offered by the server, as described in figure [Figure 4.2](#).

When data is available for an event, the server application sends the event data to communication management middleware, that notifies all subscribing client applications. The subscribers can then fetch the event samples, using `GetNewSamples()`, either directly or via a callback (defined by `SetReceiveHandler()`) that is triggered by the notification.

Triggers are used by the server to notify when a specific condition occurs. It does not transfer any data. It uses the same subscription and notification mechanisms as events.




Figure 4.2: ara::com event based communication

4.4 ara::com Method based communication

With method based communication a client application calls a method that is executed on the remote server. This is shown in figure Figure 4.3. The method may, or may not return a value to the client. If a return value is provided, the `ara::core::Future` and `ara::core::Promise` pattern is used to give a possibility of a non-blocking behavior for the communication. See chapter 5.3.6 for details.

The server can be configured for different processing modes of method invocations. The options are

- **Event-driven, concurrent (kEvent)**: Incoming service method calls are processed in an event based manner.

- **Event-driven, sequential** (`kEventSingleThread`): Same as `kEvent` on single thread basis.
- **Polling** (`KPoll`): Incoming service method calls need to be explicitly processed in polling manner by calling `ProcessNextMethodCall`.




Figure 4.3: ara::com method based communication

4.5 ara::com Field based communication

With field based communication a server can provide a value for some data that a client can access or update at any time. The functionality of a field can be viewed as a combination of event and methods:

- Like an event the client can subscribe to changes of the value. The client will be notified using the same notification mechanisms as for events (defined by `SetReceiveHandler()`).
- Using methods, the client can retrieve the value by calling a get-operation (`Get()`), or update the value by calling a set-operation for the field in the proxy (`Set()`).

On the server side, the field is handled in the skeleton implementation by

- Defining a callback that is called when the value is updated by a client (defined by `RegisterSetHandler()`).
- Calling an update-method when a new value shall be published to the clients (`Update()`).

4.6 Data Type Abstractions

`ara::com` API introduces specific data types, which are used throughout its various interfaces. They can roughly be divided into the following classes:

- Pointer types: for pointers to data transmitted via middleware
- Collection types: for collections of data transmitted via middleware.
- Types for async operation result management: `ara::com` relies on AUTOSAR AP specific data types (see [2]), which are specific versions of C++ `std::future/std::promise`
- Function wrappers: for various application side callback or handler functions to be called by the middleware

`ara::com` defines signature and expected behavior of those types, but does not provide an implementation. The idea of this approach is, that platform vendors could easily come up with their own optimized implementation of those types.

This is obvious for collection and pointer types as one of the major jobs of an `IPC` implementation has to deal with memory allocation for the data which is exchanged between middleware users.

Being able to provide their own implementations allows to optimize for their chosen memory model.

For most of the types `ara::com` provides a default mapping to existing C++ types in `ara/com/types.h`. The default mapping of the types provided by [2] can be found in, e.g. `ara/core/future.h` or `ara/core/promise.h`. This default mapping decision could be reused by an AP product vendor.

The default mapping provided by `ara::com` even has a real benefit for a product vendor, who wants to implement its own variant: He can validate the functional behavior of his own implementation against the implementation of the default mapping.

4.7 Error Handling

`ara::com` API follows the concepts of error handling described in chapter "Error handling" in [2]. Recoverable `Errors` will be returned via an `ara::core::ErrorCode` embedded into a `ara::core::Result`, which either holds a valid return value or the `ara::core::ErrorCode`.

For each function in the `ara::com` API a set of predefined `ara::core::ErrorCodes` from the error domain `ara::com::ComErrorDomain`, (or from `ara::com::e2e::E2EErrorDomain` for E2E checks) are defined. These errors should be handled by the application that is using the API. Besides these a stack vendor may also define additional error codes, that might need to be handles as well.

`Application Errors` within `ara::com` API can only occur in the context of a call of a `SI` method and is therefore fully covered in subsection [Section 5.3.6](#) and subsection [Section 5.4.6](#).

Exceptions in the `ara::com` API are only used in case of `Violations` or `Corruptions`. These are non-recoverable and should normally not be handled by the application developer.

4.8 Service Connection Approach

4.8.1 Instance Identifiers and Instance Specifiers

Instance identifiers, which get used at proxy and as well at skeleton side, are such a central concept, that their explanation is drawn here — before the detailed description of `ara::com` proxies and skeletons in upcoming chapters.

Instance identifiers are used within `ara::com`, on client/proxy side, when a specific instance of a service shall be searched for or — at the server/skeleton side — when a specific instance of a service is created.

At `ara::com` API level the instance identifier is generally a technical binding specific identifier.

Therefore the concrete content/structure of which such an instance identifier consists, is totally technology specific: So f.i. SOME/IP is using 16 bit unsigned integer identifiers to distinguish different instances of the same service type, while DDS (DDS-RPC) uses *string<256>* as `service_instance_name`.

Independant of the binding technology the abstract facade of any concrete instance identifier shall apply to this signature at `ara::com` API level in namespace `ara::com`:

```
1 class InstanceIdentifier {
2   public:
3
4   static ara::core::Result<InstanceIdentifier>
```

```
5     Create(StringView serializedFormat) noexcept;
6
7     explicit InstanceIdentifier
8         (ara::core::StringView serializedFormat);
9     ara::core::StringView ToString() const;
10    bool operator==(const InstanceIdentifier& other) const;
11    bool operator<(const InstanceIdentifier& other) const;
12    InstanceIdentifier& operator=(const InstanceIdentifier& other);
13 };
```

Listing 4.1: InstanceIdentifier class

As you can see the instance identifier interface `ara::com::InstanceIdentifier` provides a `ctor` taking a string, which means it can be constructed from a string representation. It also provides a `ToString()` method, which allows to get a stringified representation of the technology specific `ara::com::InstanceIdentifier`.

This pair of `ctor` taking a string representation and the possibility to write out the string representation makes the `ara::com::InstanceIdentifier` "serializable". This allows it to be transferred, persisted, later re-used, ... (more on potential use cases later).

Introspection into this string (trying to interpret the content) makes no sense for the user of `ara::com`. As mentioned: The content will be highly middleware product/binding specific!

Since it is a core feature, that the technical binding used by an `ara::com` based application is defined/specified by the integrator during deployment any expectations from an `ara::com` software developer regarding its content/structure are typically invalid. Logging it/tracing it out to a log channel might be helpful for debug analysis however.

Then, where does the software-developer get such a highly binding specific `ara::com::InstanceIdentifier` to be used in `ara::com` API calls?

The answer is: By an `ara::com` provided functionality, which translates a logical local name used typically by the software developer in his realm into the technology/binding specific `ara::com::InstanceIdentifier`. This indirection masters both challenges:

- developer using `ara::com` does not need to know anything about bindings and their specifics
- Integrators can adapt bindings in deployments

The local name from which the `ara::com::InstanceIdentifier` is constructed comes basically from the AUTOSAR meta-model, describing your software component model.

The requirement for this local name — we will call it "instance specifier" from now on — is, that it is unambiguous within an executable. It has basically the form:

```
/<context 0>/<context 1>/.../<context N>/<port name>
```

The C++ representation of such an "instance specifier" is the class `ara::core::InstanceSpecifier`. Structurally it looks similar to the `ara::com::InstanceIdentifier`:

```
1 class InstanceSpecifier final
2 {
3 public:
4     // ctor to build specifier from AUTOSAR short name identifier
5     // with '/' as separator between package names
6     static Result<InstanceSpecifier> Create(StringView metaModelIdentifier);
7     explicit InstanceSpecifier(StringView metaModelIdentifier);
8     InstanceSpecifier(const InstanceSpecifier& other);
9     InstanceSpecifier(InstanceSpecifier&& other) noexcept;
10    InstanceSpecifier& operator=(const InstanceSpecifier& other);
11    InstanceSpecifier& operator=(InstanceSpecifier&& other);
12    ~InstanceSpecifier() noexcept;
13    StringView ToString() const noexcept;
14    bool operator==(const InstanceSpecifier& other) const noexcept;
15    bool operator==(StringView other) const noexcept;
16    bool operator!=(const InstanceSpecifier& other) const noexcept;
17    bool operator!=(StringView other) const noexcept;
18    bool operator<(const InstanceSpecifier& other) const noexcept;
19 };
```

Listing 4.2: InstanceSpecifier class

If the unambiguousness is ensured, the integrator/deployer can assign a dedicated technical binding with its specific instance IDs to those "instance specifier" via a "manifest file", which is specifically used for a distinct instantiation/execution of the executable.

This explicitly allows, to start the same executable N times, each time with a different manifest, which maps the same `ara::core::InstanceSpecifier` differently.

Details about the `ara::com` relation to the meta-model and the nature of nested contexts can be read more detailed in [Section 7.4](#).

The API `ara::com` provides the following function, to do the translation from the `ara::core::InstanceSpecifier` (local name in the software developers realm) to the technical `ara::com::InstanceIdentifier`:


```
1 namespace ara {
2 namespace com {
3 namespace runtime {
4 ara::core::Result<ara::com::InstanceIdentifierContainer> ResolveInstanceIDs
    (ara::core::InstanceSpecifier modelName);
5 }
6 }
7 }
```

Listing 4.3: InstanceSpecifier Resolution

Why this API does return an `ara::com::InstanceIdentifierContainer`, which represents a collection of `ara::com::InstanceIdentifier` is in need of explanation: AUTOSAR supports, that the integrator may configure multiple technical bindings behind one abstract identifier visible to the software component developer.

This feature is called multi-binding and referred to at different parts in this document (you find a more detailed explanation in [Section 7.3](#)).

Using multi-binding on the skeleton/server side is a common use case, since it simply allows different clients to use their preferred binding, when contacting the server.

Contrary using multi-binding on the proxy/client side is a rather exotic one. E.g. it could be used to support some fail-over approaches (if binding A does not work, fall back on binding B).

So the possible returns for a call of `ResolveInstanceIDs()` are:

- empty list: The integrator failed to provide a mapping for the abstract identifier. This most likely is a configuration error.
- list with one element: The common case. Mapping to one concrete instance id of one concrete technical binding.
- list with more than one element: Mapping to multiple technical instances with possibly multiple technical bindings.

Technically the middleware implementation of `ResolveInstanceIDs()` does a lookup of the `ara::core::InstanceSpecifier` from the `service instance manifest` bundled within the process.

Therefore the `ara::core::InstanceSpecifier` must be unambiguous within the `bundled service instance manifest`.

4.8.2 When to use InstanceIdentifier versus InstanceSpecifier

According to the previous explanations, the impression may have arisen that a software developer always has to resolve `ara::core::InstanceSpecifier` to `ara::com::InstanceIdentifier` manually (by a call to `ResolveInstanceIDs()`) first, before using `ara::com` APIs, which need instance identifier information.

This would be indeed a bit awkward as we already mentioned, that the "typical" approach for a software developer, which implements an Adaptive AUTOSAR SWC, is to use abstract "instance specifiers" from the realm of the software component model.

As you will see in the upcoming chapters, which detail the APIs on the proxy and skeleton side, `ara::com` provides typically function overloads, which either take `ara::com::InstanceIdentifier` OR `ara::core::InstanceSpecifier`, freeing the developer in the most common use cases, where he simply uses `ara::core::InstanceSpecifier` from explicitly calling `ResolveInstanceIDs()`.

This means, that the direct use of `ara::com::InstanceIdentifier` and manual resolution of `ara::core::InstanceSpecifier` is intended more for power users with rather specific/exotic use cases. Some examples will be given in the chapters, where the corresponding `ara::com` API overrides at proxy/skeleton side are discussed.

The fundamental difference between the two variants is this: An `ara::com::InstanceIdentifier` can be exchanged more easily between Adaptive Applications/processes!

As they already exactly contain all the technology specific information and do not need any further resolution via content of a `service instance manifest` such a serialized `ara::com::InstanceIdentifier` can be reconstructed within a different process and be used as long as his process has access to the same binding technology the `ara::com::InstanceIdentifier` is based upon.

4.8.2.1 Transfer of an InstanceIdentifier

As discussed before the `ara::com::InstanceIdentifier` should only be used for "power users" since its format is stack vendor dependent and it contains technology binding information. The transfer or the storage of an `ara::com::InstanceIdentifier` may be very risky, therefore. As the transfer binding may not exist anymore after the transfer or re-storing or the `ara::com::InstanceIdentifier` of stack vendor A may be interpreted by an application using the stack of vendor B.

5 Detailed API description

5.1 High Level API Structure

`ara::com` provides an API that supports the AUTOSAR service model. The services have methods, events, fields and triggers.

- **Methods:** Execute a function in the Service Application which can also return a value (e.g. Calibrate method).
- **Events:** The Service Application sends an event (may also include a value) when specific conditions occur (e.g. Brake event). A Client Application can subscribe to events.
- **Fields:** Have a value at any time, like a status value. Can be read using Get or modified using Set (e.g. UpdateRate field). A Client Application can be notified when a Field value changes.
- **Triggers:** The Service Application sends a trigger when specific conditions occur. A Client Application can subscribe to triggers.

As described in [Section 4.1](#), Client and Service Application communicate with each other and therefore the API supports methods, events and fields in both sides. This means that the API defines interfaces for sending and receiving events, provides and calls service methods, register handlers for field setters and getters amongst others.

The `ara::com` API also defines ctors/dtors to create and destroy instances for Proxy and Skeleton classes.

Finally, the `ara::com` API also provides methods to offer / find services and subscribe / unsubscribe to events.

5.2 API Elements

The following subchapters will guide through the different API elements, which `ara::com` defines. Since we will give code examples for various artifacts and provide sample code how to use those APIs from a developer perspective, it is a good idea to have some uniformity in our examples.

So we will use a virtual service (interface) called "RadarService". The following is a kind of a semi-formal description, which should give you an impression of what this "RadarService" provides/does and might be easier to read than a formal AUTOSAR ARXML service description:

```
1 RadarService {
2   // types used within service
3   type RadarObjects {
4     active : bool
5     objects : array {
6       elementtype: uint8
```

```

7     size: variable
8   }
9 }
10
11 type Position {
12   x: uint32
13   y: uint32
14   z: uint32
15 }
16
17 // events provided by service
18 event BrakeEvent {
19   type: RadarObjects
20 }
21
22 // fields provided by service
23 field UpdateRate {
24   type: uint32
25   get: true
26   set: true
27 }
28
29 error CalibrationFailed {
30   errorCode : 1
31   errorContext {
32     failureText : string
33   }
34 }
35
36 error InvalidConfigString {
37   errorCode : 2
38   errorContext {
39     invalidConfig : string
40     currentValidConfig : string
41   }
42 }
43
44 // methods provided by service
45 method Calibrate {
46   param configuration {
47     type: string
48     direction: in
49   }
50   param result {
51     type: bool
52     direction: out
53   }
54   raises {
55     CalibrationFailed
56     InvalidConfigString
57   }
58 }
59
60 method Adjust {
61   param target_position {
62     type: Position

```

```

63     direction: in
64   }
65   param success {
66     type: bool
67     direction: out
68   }
69   param effective_position {
70     type: Position
71     direction: out
72   }
73 }
74
75 oneway method LogCurrentState {}
76 }
    
```

Listing 5.1: RadarService Definition

So the example service `RadarService` provides an event “BrakeEvent”, which consists of a structure containing a flag and a variable length array of `uint8` (as extra payload).

Then it provides a field “UpdateRate”, which is of `uint32` type and supports `get` and `set` calls and finally it provides three methods.

Method “Adjust”, to position the radar, which contains a target position as in-parameter and two out-parameters. One to signal the success of the positioning and one to report the final (maybe deviating) effective position.

The method “Calibrate” to calibrate the radar, getting an configuration string as in-parameter and returning a success indicator as out-parameter. This method may raise two different application errors, in case the calibration failed: “CalibrationFailed” and “InvalidConfigString”.

The method “LogCurrentState” is a one way method, which means, that no feedback is returned to the caller, if the method is executed at all and with which outcome. It instructs the service `RadarService` to output its current state into its local log files.

5.3 Proxy Class

The Proxy class is generated from the [SI](#) description of the AUTOSAR meta model.

`ara::com` does standardize the interface of the generated Proxy class. The toolchain of an AP product vendor will generate a Proxy implementation class exactly implementing this interface.

Note: Since the interfaces the Proxy class has to provide are defined by `ara::com`, a generic (product independent) generator could generate an abstract class or a mock class against which the application developer could implement his service consumer application. This perfectly suits the platform vendor independent development of Adaptive AUTOSAR SWCs.

ara::com expects proxy related artifacts inside a namespace "proxy". This namespace is typically included in a namespace hierarchy deduced from the service definition and its context.

5.3.1 Proxy Class API's

- FindService()
- StartFindService()
- StopFindService()
- Subscribe()
- Unsubscribe()
- GetSubscriptionState()
- SetSubscriptionStateChangeHandler()
- UnsetSubscriptionStateChangeHandler()
- GetNewSamples()
- GetFreeSampleCount()
- SetReceiveHandler()
- UnsetReceiveHandler()
- ResolveInstanceIDs()
- Field::Get()
- Field::Set()

5.3.2 RadarService Proxy Class Example

```

1 class RadarServiceProxy {
2     public:
3         /**
4          * \brief Implementation is platform vendor specific
5          *
6          * A HandleType must contain the information that is needed to create
7          * a proxy.
8          *
9          * This information shall be hidden.
10         * Since the platform vendor is responsible for creation of handles,
11         the
12         * ctor signature is not given as it is not of interest to the user.
13         */
14         class HandleType {
15             /**

```

```

15     * \brief Two ServiceHandles are considered equal if they represent
16     * the same service instance.
17     *
18     * \param other
19     *
20     * \return bool
21     */
22     inline bool operator==(const HandleType &other) const;
23     const ara::com::InstanceIdentifier &GetInstanceId() const;
24 };
25
26 /**
27  * StartFindService does not need an explicit version parameter as this
28  * is internally available in ProxyClass.
29  * That means only compatible services are returned.
30  *
31  * \param handler this handler gets called any time the service
32  * availability of the services matching the given
33  * instance criteria changes. If you use this variant of
34  * FindService, the Communication Management has to
35  * continuously monitor the availability of the services
36  * and call the handler on any change.
37  *
38  * \param instanceId which instance of the service type defined
39  * by T shall be searched/found.
40  *
41  * \return a handle for this search/find request, which shall
42  * be used to stop the availability monitoring and related
43  * firing of the given handler. (\see StopFindService())
44  */
45     static ara::core::Result<ara::com::FindServiceHandle> StartFindService(
46     ara::com::FindServiceHandler<RadarServiceProxy::HandleType> handler,
47     ara::com::InstanceIdentifier instanceId);
48
49     /**
50     * This is an overload of the StartFindService method using an
51     * instance specifier, which gets resolved via service instance
52     * manifest.
53     * \param instanceSpec instance specifier
54     */
55     static ara::core::Result<ara::com::FindServiceHandle> StartFindService
56     (
57     ara::com::FindServiceHandler<RadarServiceProxy::HandleType> handler,
58     ara::core::InstanceSpecifier instanceSpec);
59
60     /**
61     * Method to stop finding service request (see above)
62     */
63     static void StopFindService(ara::com::FindServiceHandle handle);
64
65     /**
66     * Opposed to StartFindService(handler, instance) this version
67     * is a "one-shot" find request, which is:
68     * - synchronous, i.e. it returns after the find has been done
69     * and a result list of matching service instances is
70     * available. (list may be empty, if no matching service

```

```

70     *   instances currently exist)
71     * - does reflect the availability at the time of the method
72     *   call. No further (background) checks of availability are
73     *   done.
74     *
75     * \param instanceId which instance of the service type defined
76     * by T shall be searched/found.
77     *
78     */
79     static ara::core::Result<ara::com::ServiceHandleContainer
80         <RadarServiceProxy::HandleType>>
81         FindService(ara::com::InstanceIdentifier instanceId);
82
83     /**
84     * This is an overload of the FindService method using an
85     * instance specifier, which gets resolved via service instance
86     * manifest.
87     */
88     static ara::core::Result<ara::com::ServiceHandleContainer
89         <RadarServiceProxy::HandleType>>
90         FindService(ara::core::InstanceSpecifier instanceSpec);
91
92     /**
93     * \brief The proxy can only be created using a specific
94     * handle which identifies a service.
95     *
96     * This handle can be a known value which is defined at
97     * deployment or it can be obtained using the
98     * ProxyClass::FindService method.
99     *
100    * \param handle The identification of the service the
101    * proxy should represent.
102    */
103    explicit RadarServiceProxy(HandleType &handle);
104
105    /**
106    * proxy instances are not copy constructible.
107    */
108    RadarServiceProxy(RadarServiceProxy &other) = delete;
109
110    /**
111    * proxy instances are not copy assignable
112    */
113    RadarServiceProxy& operator=(const RadarServiceProxy &other) = delete;
114
115    /**
116    * \brief Public member for the BrakeEvent
117    */
118    events::BrakeEvent BrakeEvent;
119
120    /**
121    * \brief Public Field for UpdateRate
122    */
123    fields::UpdateRate UpdateRate;
124
125    /**

```



```
126     * \brief Public member for the Calibrate method
127     */
128     methods::Calibrate Calibrate;
129
130     /**
131     * \brief Public member for the Adjust method
132     */
133     methods::Adjust Adjust;
134
135     /**
136     * \brief Public member for the LogCurrentState fire-and-forget method
137     */
138     methods::LogCurrentState LogCurrentState;
139 };
```

Listing 5.2: RadarService Proxy

5.3.3 Constructor and Handle Concept

As you can see in the Listing [Listing 5.2](#) `ara::com` prescribes the Proxy class to provide a constructor. This means, that the developer is responsible for creating a proxy instance to communicate with a possibly remote service.

The `ctor` takes a parameter of type `ProxyHandleType` — an inner class of the generated proxy class. Probably the immediate question then is: *"What is this handle and how to create it/where to get it from?"*

What it is, should be straightforward: After the call to the `ctor` you have a proxy instance, which allows you to communicate with the service, therefore the handle has to contain the needed addressing information, so that the Communication Management binding implementation is able to contact the service.

What exactly this address information contains is totally dependent on the binding implementation/technical transport layer!

That already partly answers the question *"how to create/where to get it"*: Really creating is not possible for an application developer as he is — according to AUTOSAR core concepts — implementing his application AP product and therefore Communication Management independent.

The solution is, that `ara::com` provides the application developer with an API to find service instances, which returns such handles.

This part of the API is described in detail here: [Section 5.3.4](#). The co-benefit from this approach — that proxy instances can only be created from handles, which are the result of a "FindService" API — is, that you are only able to create proxies, which are really backed by an existing service instance.

So the question which probably might come up here: Why this indirection, that an application developer first has to call some `ara::com` provided functionality, to get a handle, which I then have to use in a `ctor` call? `ara::com` could have given back directly a proxy instance instead of a handle from "FindService" functionality.

The reason for that could be better understood, after reading how `ara::com` handles the access to events ([Section 5.3.5](#)). But what is sufficient to say at this point is, that a proxy instance contains certain state.

And because of this there are use cases, where the application developer wants to use different instances of a proxy, all "connected" to the same service instance.

So if you just accept, that there are such cases, the decision for this indirection via handles becomes clear: `ara::com` cannot know, whether an application developer wants always the same proxy instance (explicitly sharing state) or always a new instance each time he triggers some "FindService" functionality, which returns a proxy for exactly the same service instance.

So by providing this indirection/decoupling the decision is in the hands of the `ara::com` user.

Instances of the Proxy class on the other hand are neither copy constructible nor copy assignable! This is an explicit design decision, which complements the idea of forcing the construction via `HandleType`.

The instances of a proxy class might be very resource intensive because of owning event/field caches, registered handlers, complex state,... and so on. Thus, when allowing copy construction/copy assignment, there is a risk that such copies are done unintended.

So — in a nutshell — forcing the user to go the route via `HandleType` for Proxy creation shall sensitize him, that this decision shall be well thought out.

5.3.4 Finding Services

The Proxy class provides class (static) methods to find service instances, which are compatible with the Proxy class.

Since the availability of service instances is dynamic by nature, as they have a life cycle, `ara::com` provides two different ways to do a "FindService" for convenience in general:

- `StartFindService` is a class method, which starts a continuous "FindService" activity in the background, which notifies the caller via a given callback anytime the availability of instances of the service changes.
- `FindService` is a one-off call, which returns available instances at the point in time of the call.

Both of those methods come in two different overrides, depending on the instance identifier approach taken (see [Section 4.8.1](#)):

- one taking an `ara::com::InstanceIdentifier`
- one taking an `ara::core::InstanceSpecifier`

Note that only technical bindings will be used for finding/searching, which are configured for the corresponding [SI](#) within the service instance manifest in the form of a [SI](#) deployment.

The synchronous one-off variant [FindService](#) returns a container of handles (see [Section 5.3.3](#)) for the matching service instances, which might also be empty, if no matching service instance is currently available.

Opposed to that, the [StartFindService](#) returns a [FindServiceHandle](#), which can be used to stop the ongoing background activity of monitoring service instance availability via call to [StopFindService](#).

The first (and specific for this variant) parameter to [StartFindService](#) is a user provided handler function with the following signature:

```
1 using FindServiceHandler = std::function<void(ServiceHandleContainer<T  
>, FindServiceHandle)>;
```

Any time the binding detects, that the availability of service instances matching the given instance criteria in the call to [StartFindService](#) has changed, it will call the user provided handler with an updated list of handles of the now available service instances.

Right after being called, [StartFindService](#) behaves similar to [FindService](#) in the sense, that it will fire the user provided handler function with the currently available service instances, which might be also an empty handle list.

After that initial callback, it will call the provided handler again in case of changes of this initial service availability.

Note, that it is explicitly allowed, that the [ara::com](#) user/developer does call [StopFindService](#) within the user provided handler.

For this purpose, the handler explicitly gets the [FindServiceHandle](#) argument. The handler needs not to be re-entrant. This means, that the binding implementer has to care for serializing calls to the user provided handler function.

Note, that [ServiceHandleContainer](#) can be implemented as an allocating or non-allocating container, when used either as a return value of [FindService](#) or as a parameter to [FindServiceHandler](#), as long as it fulfils general and sequence container requirements of the C++ programming language.

5.3.4.1 Auto Update Proxy instance

Regardless whether you use the one-off [FindService](#) or the [StartFindService](#) variant, in both cases you get a handle identifying the — possibly remote — service instance, from which you then create your proxy instance.

But what happens if the service instance goes down and later comes up again e.g. due to some life cycle state changes? Can the existing proxy instance at the service consumer side still be re-used later, when the service instance gets available again?

The good news is: The `ara::com` design team decided to require this re-use possibility from the binding implementation as it eases the typical task of implementing service consumers.

In the service based communication universe it is expected, that during the life time of the entire system (e.g. vehicle) service provider and consumer instances are starting up and going down again due to their own life cycle concepts frequently.

To deal with that, there is the service discovery infrastructure, where the life cycle of service providers and consumers is monitored in terms of service offerings and service (re)subscriptions!

If a service consumer application has instantiated a service proxy instance from a handle returned from some of the `Find Service` variants, the sequence which might possibly occur is shown in the figure below.

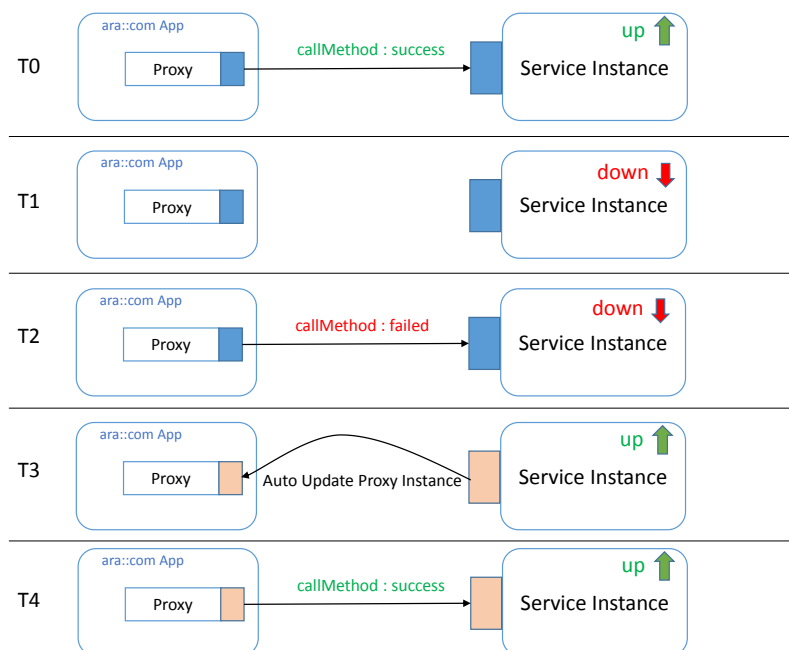


Figure 5.1: Auto Updating of Proxy Instance

Explanation of figure 5.1:

- **T0:** The service consumer may successfully call a service method of that proxy (and `GetSubscriptionState()` on subscribed events will return `kSubscribed` according to 5.3.5.2).
- **T1:** The service instance goes down, correctly notified via service discovery.
- **T2:** A call of a service method on that proxy will lead to a `kServiceNotAvailable` error, since the targeted service instance of the call does not exist anymore.

Correspondingly `GetSubscriptionState()` on any subscribed event will return `kSubscriptionPending` (see also 5.3.5.2) at this point even if the event has been successfully subscribed (`kSubscribed`) before.

- **T3:** The service instance comes up again, notified via service discovery infrastructure. The Communication Management at the proxy side will be notified and will silently update the proxy object instance with a possibly changed transport layer addressing information. This is illustrated in the figure with transport layer part of the proxy, which changed the color from blue to rose.
- **T4:** Consequently service method calls on that proxy instance will succeed again and `GetSubscriptionState()` on events which the service consumer had subscribed before, will return `kSubscribed` again.

This convenience behavior of a proxy instance saves the implementer of a service consumer from either:

- polling via `GetSubscriptionState()` on events, which indicates that service instance has gone down
- re-triggering a one-off `FindService` to get a new handle.

or:

- registering a `FindServiceHandler`, which gets called in case service instance gets down or up with a new handle.

and then to recreate a proxy instance from the new handle (and redo needed event subscribe calls).

Note, in case you have registered a `FindServiceHandler`, then the binding implementation must assure, that it does the “auto updating” of existing proxy instances **before** it calls the registered `FindServiceHandler`!

The reason for this is: It shall be supported, that the application developer can interact successfully with an existing proxy instance within the `FindServiceHandler`, when the handle of the proxy instance is given in the call, signaling, that the service instance is up again.

This expectation is shown in the following code snippet:

```

1 /**
2  * Reference to radar instance, we work with,
3  * initialized during startup
4  */
5 RadarServiceProxy *myRadarProxy;
6
7 void radarServiceAvailabilityHandler(ServiceHandleContainer<
8   RadarServiceProxy::HandleType> curHandles, FindServiceHandle handle) {
9   for (RadarServiceProxy::HandleType handle : curHandles) {
10    if (handle.GetInstanceId() == myRadarProxy->GetHandle().
11    GetInstanceId()) {
12      /**
13       * This call on the proxy instance shall NOT lead to an
14       * exception,
15       * regarding service instance not reachable, since proxy
16       * instance
17       * should be already auto updated at this point in time.
18       */
19      ara::core::Future<Calibrate::Output> out =
20      myRadarProxy->Calibrate("test");
21
22      // ... do something with out.
23    }
24  }
25 }

```

Listing 5.3: Access to proxy instance within FindService handler

5.3.5 Events

For each event the remote service provides, the proxy class contains a member of a event specific wrapper class. In our example the member has the name `BrakeEvent` and is of type `events::BrakeEvent`.

As you see in 5.2 all the event classes needed for the proxy class are generated inside a specific namespace `events`, which is contained inside the `proxy` namespace.

The member in the proxy is used to access events/event data, which are sent by the service instance our proxy is connected to. Let’s have a look at the generated event class for our example:

```

1 class BrakeEvent {

```

```

2     /**
3     * \brief Shortcut for the events data type.
4     */
5     using SampleType = RadarObjects;
6
7     /**
8     * \brief The application expects the CM to subscribe the event.
9     *
10    * The Communication Management shall try to subscribe and resubscribe
11    * until \see Unsubscribe() is called explicitly.
12    * The error handling shall be kept within the Communication Management
13    *
14    * The function returns immediately. If the user wants to get notified,
15    * when subscription has succeeded, he needs to register a handler
16    * via \see SetSubscriptionStateChangeHandler(). This handler gets
17    * then called after subscription was successful.
18    *
19    * \param maxSampleCount maximum number of samples, which can be held.
20    */
21    ara::core::Result<void> Subscribe(size_t maxSampleCount);
22
23    /**
24    * \brief Query current subscription state.
25    *
26    * \return Current state of the subscription.
27    */
28    ara::com::SubscriptionState GetSubscriptionState() const;
29
30    /**
31    * \brief Unsubscribe from the service.
32    */
33    void Unsubscribe();
34
35    /**
36    * \brief Get the number of currently free/available sample slots.
37    *
38    * \return number from 0 - N (N = count given in call to Subscribe())
39    *         or an ErrorCode in case of number of currently held samples
40    *         already exceeds the max number given in Subscribe().
41    */
42    size_t GetFreeSampleCount() const noexcept;
43
44    /**
45    * Setting a receive handler signals the Communication Management
46    * implementation to use event style mode.
47    * I.e. the registered handler gets called asynchronously by the
48    * Communication Management as soon as new event data arrives for
49    * that event. If the user wants to have strict polling behavior,
50    * where no handler is called, NO handler should be registered.
51    *
52    * Handler may be overwritten anytime during runtime.
53    *
54    * Provided Handler needs not to be re-entrant since the
55    * Communication Management implementation has to serialize calls
56    * to the handler: Handler gets called once by the MW, when new
    
```

```
57     * events arrived since the last call to GetNewSamples().
58     *
59     * When application calls GetNewSamples() again in the context of the
60     * receive handler, MW must - in case new events arrived in the
61     * meantime - defer next call to receive handler until after
62     * the previous call to receive handler has been completed.
63     */
64     ara::core::Result<void> SetReceiveHandler(ara::com::EventReceiveHandler
        handler);
65
66     /**
67     * Remove handler set by SetReceiveHandler()
68     */
69     ara::core::Result<void> UnsetReceiveHandler();
70
71     /**
72     * Setting a subscription state change handler, which shall get
73     * called by the Communication Management implementation as soon
74     * as the subscription state of this event has changed.
75     *
76     * Communication Management implementation will serialize calls
77     * to the registered handler. If multiple changes of the
78     * subscription state take place during the runtime of a
79     * previous call to a handler, the Communication Management
80     * aggregates all changes to one call with the last/effective
81     * state.
82     *
83     * Handler may be overwritten during runtime.
84     */
85     ara::core::Result<void> SetSubscriptionStateChangeHandler(
86         ara::com::SubscriptionStateChangeHandler handler);
87
88     /**
89     * Remove handler set by SetSubscriptionStateChangeHandler()
90     */
91     void UnsetSubscriptionStateChangeHandler();
92
93     /**
94     * \brief Get new data from the Communication Management
95     * buffers and provide it in callbacks to the given callable f.
96     *
97     * \pre BrakeEvent::Subscribe has been called before
98     * (and not be withdrawn by BrakeEvent::Unsubscribe)
99     *
100    * \param f
101    * \parblock
102    * callback, which shall be called with new sample.
103    *
104    * This callable has to fulfill signature
105    * void(ara::com::SamplePtr<SampleType const>)
106    * \parblockend
107    *
108    * \param maxNumberOfSamples
109    * \parblock
110    * upper bound of samples to be fetched from middleware buffers.
111    * Default value means "no restriction", i.e. all newly arrived samples
```



```

112     * are fetched as long as there are free sample slots.
113     * \parblockend
114     *
115     * \return Result, which contains the number of samples,
116     * which have been fetched and presented to user via calls to f or an
117     * ErrorCode in case of error (e.g. precondition not fulfilled)
118     */
119     template <typename F>
120     ara::core::Result<size_t> GetNewSamples(
121         F&& f,
122         size_t maxNumberOfSamples = std::numeric_limits<size_t>::max());
123 };
    
```

Listing 5.4: Proxy side BrakeEvent Class

The data type of the event data in our example event is `RadarObjects` (see [Listing 5.1](#)). The first you encounter is the using-directive which assigns the generic name `SampleType` to the concrete type, which is then used throughout the interface.

5.3.5.1 Event Subscription and Local Cache

The mere fact, that there exists a member of the event wrapper class inside the proxy instance does not mean, that the user gets instant access to events raised/sent out by service instance.

First you have to “subscribe” for the event, in order to tell the Communication Management, that you are now interested in receiving events.

For that purpose the event wrapper class of `ara::com` provides the method

```

1     /**
2     * \brief The application expects the CM to subscribe the event.
3     *
4     * ....
5     *
6     * \param maxSampleCount maximum number of samples, which can be held.
7     */
8     ara::core::Result<void> Subscribe(size_t maxSampleCount);
    
```

This method expects a parameter `maxSampleCount`, which basically informs Communication Management implementation, how many event samples the application intends to hold at maximum. Therefore — with calling this method, you not only tell the Communication Management, that you now are interested in receiving event updates, but you are at the same time setting up a “local cache” for those events bound to the event wrapper instance with the given `maxSampleCount`.

This cache is allocated and filled by the Communication Management implementation, which hands out smartpointers to the application for accessing the event sample data. How that works in detail is described in [Section 5.3.5.3](#)).

5.3.5.2 Monitoring Event Subscription

The call to the `Subscribe()` method is asynchronous by nature. This means that at the point in time `Subscribe()` returns, it is just the indication, that the Communication Management has accepted the order to care for subscription.

The subscription process itself may (most likely, but depends on the underlying `IPC` implementation) involve the event provider side. Contacting the possibly remote service for setting up the subscription might take some time.

So the binding implementation of the subscribe is allowed to return immediately after accepting the subscribe, even if for instance the remote service instance has not yet acknowledged the subscription (in case the underlying `IPC` would support mechanism like acknowledgment at all). If the user — after having called `Subscribe()` — wants to get feedback about the success of the subscription, he might call:

```
1  /**
2   * \brief query current subscription state.
3   *
4   * \return current state of the subscription.
5   */
6  ara::com::SubscriptionState GetSubscriptionState() const;
```

In the case the underlying `IPC` implementation uses some mechanism like a subscription acknowledge from the service side, then an immediate call to `GetSubscriptionState()` after `Subscribe()` may return `kSubscriptionPending`, if the acknowledge has not yet arrived.

Otherwise — in case the underlying `IPC` implementation gets instant feedback, which is very likely for local communication — the call might also already return `kSubscribed`.

If the user needs to monitor the subscription state, he has two possibilities:

- Polling via `GetSubscriptionState()`
- Registering a handler, which gets called, when the subscription state changes

The first possibility by using `GetSubscriptionState()` we have already described above. The second possibility relies on using the following method on the event wrapper instance:

```
1  /**
2   * Setting a subscription state change handler, which shall get called
3   * by the Communication Management implementation as soon as the
4   * subscription state of this event has changed.
5   *
6   * Handler may be overwritten during runtime.
7   */
8  ara::core::Result<void> SetSubscriptionStateChangeHandler
9    (ara::com::SubscriptionStateChangeHandler handler);
```

Here the user may register a handler function, which has to fulfill the following signature:

```
1  enum class SubscriptionState { kSubscribed, kNotSubscribed,  
   kSubscriptionPending };  
2  using SubscriptionStateChangeHandler = std::function<void(  
   SubscriptionState)>;
```

Anytime the subscription state changes, the Communication Management implementation calls the registered handler. A typical usage pattern for an application developer, who wants to get notified about latest subscription state, would be to register a handler **before** the first call to `Subscribe()`.

After having accepted the “subscribe order” the Communication Management implementation will call the handler first with argument `SubscriptionState.kSubscriptionPending` and later — as it gets acknowledgment from the service side — it will call the handler with argument `SubscriptionState.kSubscribed`.

Again the note: If the underlying implementation does not support a subscribe acknowledgment from the service side, the implementation could also skip the first call to the handler with argument `SubscriptionState.kSubscriptionPending` and **directly** call it with argument `SubscriptionState.kSubscribed`.

Calls to the registered “subscription state change” handler are done fully asynchronous. That means, they can even happen, while the call to `Subscribe()` has not yet returned. The user has to be aware of this!

Once the user has registered such a “subscription state change” handler for a certain event, he may receive multiple calls to this handler. Not only initially, when the state changes from `SubscriptionState.kNotSubscribed` to `SubscriptionState.kSubscribed` (eventually via an intermediate step `SubscriptionState.kSubscriptionPending`), but also anytime later as the service providing this event may have a certain life-cycle (maybe bound to certain vehicle modes).

The service might therefore toggle between availability and (temporarily) unavailability or it might even unexpectedly crash and restart. Those changes of the availability of the service instance providing the event may be visible to the proxy side Communication Management implementation.

The Communication Management therefore will fire the registered “subscription state change” handler, whenever it detects such changes, which have influence on the event subscription state.

Additionally (and maybe even more important) — the Communication Management implementation takes care of renewing/updating event subscriptions done by the user, whenever needed.

This mechanism is closely coupled with the “Auto Update Proxy instance” mechanism already described above (5.3.4.1): Since the Communication Management implementation monitors the availability of the service instances, the service proxies are connected to it automatically once the service is available.

The mechanism does not only “auto-update” its proxies if needed, but also “silently” re-subscribes any event subscription already done by the user, after it has updated a proxy instance.

This can be roughly seen as a very useful comfort feature — without this “re-subscribe after update”, the “auto-update” alone seemed to be a halfhearted approach.

With registration of a “subscription state change” handler, the user has now another possibility to monitor the current availability of a service! Beside the possibility to register a `FindServiceHandler` as described in 5.3.4, the user, who has registered a “subscription state change” handler, can monitor the service availability indirectly by calls to his handler.

In case the service instance, the proxy is connected to, goes down, the Communication Management calls the handler with argument `SubscriptionState.kSubscriptionPending`. As soon as the “re-subscribe after update” was successful, the Communication Management calls the handler with argument `SubscriptionState.kSubscribed`.

An `ara::com` compliant Communication Management implementation has to serialize calls to the user registered handler. I.e.: If a new subscription state change happens, while the user provided handler from a previous call of a state change is still running, the Communication Management implementation has to postpone the next call until the previous has returned.

Several subscription state changes, which happen during the runtime of a user registered state change handler, shall be aggregated to one call to the user registered handler with the effective/last state.

5.3.5.3 Accessing Event Data — aka Samples

So, after you successfully subscribed to an event according to the previous chapters, how is the access to received event data samples achieved? The event data, which is sent from the event emitter (service provider) to subscribing proxy instances is — in typical `IPC` implementations — accumulated/queued in some buffers (e.g. kernel buffers, special `IPC` implementation controlled shared memory regions, ...). So there has to be taken an **explicit** action, to get/fetch those event samples from those buffers, eventually deserialize it and then put them into the event wrapper class instance specific cache in form of a correct `SampleType`. The API to trigger this action is `GetNewSamples()`.

```

1  /**
2   * \brief Get new data from the Communication Management
3   * buffers and provide it in callbacks to the given callable f.
4   *
5   * ....
6   */
7  template <typename F>
8  ara::core::Result<size_t> GetNewSamples(
9      F&& f,
10     size_t maxNumberOfSamples = std::numeric_limits<size_t>::max());

```

As you can see, the API is a function template, due to the fact, that the first parameter `f` is a very flexible user provided `Callable`, which has to fulfill the following signature requirement: `void(ara::com::SamplePtr<SampleType const>)`.

The second argument of type `size_t` controls the maximum number of event samples, that shall be fetched/deserialized from the middleware buffers and then presented to the application in form of a call to `f`.

On a call to `GetNewSamples()`, the `ara::com` implementation checks first, whether the number of event samples held by the application already exceeds the maximum number, which it had committed in the previous call to `Subscribe()`. If so, an `ara::core::ErrorCode` is returned. Otherwise `ara::com` implementation checks, whether underlying buffers contain a new event sample and — if it's the case — deserializes it into a sample slot and then calls the application provided `f` with a `SamplePtr` pointing to this new event sample. This processing (checking for further samples in the buffer and calling back the application provided callback `f`) is repeated until either:

- there aren't any new samples in the buffers
- there are further samples in the buffers, but the application provided `maxNumberOfSamples` argument in call to `GetNewSamples()` has been reached.
- there are further samples in the buffers, but the application already exceeds its `maxSampleCount`, which it had committed in `Subscribe()`.

Within the implementation of callback `f`, which the application/user provides, it can be decided, what to do with the passed `SamplePtr` argument (i.e. by eventually doing a deep inspection of the event data): Shall the new sample be "thrown away", because it is not of interest or shall it be kept for later. To get an idea, what keeping/throwing away

of event samples means, the semantics of the `SamplePtr`, which is the access/entry point to the event sample data has to be fully understood.

The following chapter shall clarify this.

The returned `ara::core::Result` contains either an `ErrorCode` or — in the success case — the number of calls to `f`, which have been done in the context of the `GetNewSamples()` call.

5.3.5.4 Event Sample Management via `SamplePtrs`

A `SamplePtr`, which is handed over from the `ara::com` implementation to application/user layer is — from a semantical perspective — a unique-pointer (very similar to a `std::unique_ptr`): When the `ara::com` implementation hands it over an ownership transfer takes place. From now on the application/user is responsible for the lifetime management of the underlying sample. As long as the user doesn't free the sample by destroying the `SamplePtr` or by calling explicit assignment-ops/modifiers on the `SamplePtr` instance, the `ara::com` implementation can not reclaim the memory slot occupied by this sample.

Those memory-slots, in which the event sample data reside, are allocated by the `ara::com` implementation. This typically takes place in the context of the call to `Subscribe()`, where the user/application defines by parameter `maxSampleCount`, what maximum number of event data samples it wants to have concurrently accessible. Within later `GetNewSamples()` calls, the `ara::com` implementation then populates/-fills such a "sample slot" (if one is free) and passes a `SamplePtr` pointing to it in the user/application callback `f`.

In the callback implementation the user/application decides then, what to do with this passed in `SamplePtr`. If it wants to keep the sample for later access (i.e. after the return of the callback, it will make a copy at some outer scope location, where it fits in its software component architecture. The decision, whether to copy it (i.e. keep it) might simply depend on the properties/values of the event sample data. In this case the callback implementation is basically applying a "filter" on the received event samples. Since we stated, that the `SamplePtr` behaves like a `std::unique_ptr`, the above statement has to be slightly corrected: The implementation — when deciding to keep that event sample — is obviously not copying that passed in `SamplePtr`, but moving it to a outer scope location.

The small example in [Listing 5.5](#) shows — beside other things — in method `handleBrakeEventReception()` how such a callback implementation could realize simple filtering and moving of samples to a global storage with a "LastN" semantic for later use/processing.

5.3.5.5 Event-Driven vs Polling-Based access

As already promised, we fully support event-driven and polling approaches to access new data. For the polling approach no other APIs are needed than those, which we have discussed up to this point. The typical use case is, that you have an application, which is cyclically triggered to do some processing and provide its output to certain deadlines. This is the typical pattern of a regulator/control algorithm — the cyclic activation might additionally be driven by a real-time timer, which assures a minimal jitter.

In such a setup you call `GetNewSamples()` in each activation cycle and then use those updated cache data as input for the current processing iteration. Here it is fully sufficient to get the latest data to process at the time the processing algorithm is scheduled.

It would be counterproductive, if the Communication Management would notify your application anytime new data is available: This would just mean unnecessary context switches to your application process, since at the time you get the notification you do not want to process that new data as it is not time for it.

However, there are other use cases as well. If your application does not have such a cyclical, deadline driven approach, but shall simply react in case certain events occur, then setting up cyclical alarms and poll for new events via calls to `GetNewSamples()` is a bit off and vastly inefficient.

In this case you explicitly want the Communication Management to notify your application thereby issuing asynchronous context switches to your application process. We do support this flavor with the following API mechanism:

```
1   ara::core::Result<void> SetReceiveHandler(ara::com::EventReceiveHandler  
    handler);
```

This API allows you to register a user defined callback, which the Communication Management has to call in case new event data is available since the last call to `GetNewSamples()`. The registered function needs NOT to be re-entrant as the Communication Management has to serialize calls to the registered callback.

It is explicitly allowed to call `GetNewSamples()` from within the registered callback!

Note, that the user can alter the behavior between event-driven and polling style anytime as he also has the possibility to withdraw the user specific “receive handler” with the `UnsetReceiveHandler()` method provided by the event wrapper.

The following short code snippet is a simple example of how to work with events on proxy/client side. In this sample a proxy instance of type `RadarService` is created within `main` and a reception handler is registered, which gets called by the `ara::com` implementation any time new `BrakeEvent` events get received. This means, that in this example we are using the "Event-Driven" approach.

In our sample receive handler, we update our local cache with newly received events, thereby filtering out all BrakeEvent events, which do not fulfill a certain property. Afterwards we call a processing function, which processes the samples, we have decided to keep.

```
1 #include "RadarServiceProxy.hpp"
2 #include <memory>
3 #include <deque>
4
5 using namespace com::mycompany::division::radarservice;
6 using namespace ara::com;
7
8 /**
9  * our radar proxy - initially the unique ptr is invalid.
10 */
11 std::unique_ptr<proxy::RadarServiceProxy> myRadarProxy;
12
13 /**
14  * a storage for BrakeEvent samples in fifo style
15 */
16 std::deque<SamplePtr<const proxy::events::BrakeEvent::SampleType>>
    lastNActiveSamples;
17
18 /**
19  * \brief application function, which processes current set of BrakeEvent
20  * samples.
21  * \param samples
22  */
23 void processLastBrakeEvents(
24     std::deque<SamplePtr<const proxy::events::BrakeEvent::SampleType>>&
    samples) {
25     // do whatever with those BrakeEvent samples ...
26 }
27
28 /**
29  * \brief event reception handler for BrakeEvent events, which we register
    to get informed about new events.
30 */
31 void handleBrakeEventReception() {
32     /**
33      * we get newly arrived BrakeEvent events into our process space.
34      * For each sample we get passed in, we check for a certain property
35      * "active" and if it fulfills the check, we move it into our Last10-
    storage.
36      * So this few lines basically implement filtering and a LastN policy.
37      */
38     myRadarProxy->BrakeEvent.GetNewSamples(
39         [(SamplePtr<proxy::events::BrakeEvent::SampleType> samplePtr) {
40             if(samplePtr->active) {
41                 lastNActiveSamples.push_back(std::move(samplePtr));
42                 if (lastNActiveSamples.size() > 10)
43                     lastNActiveSamples.pop_front();
44             }
45         }]);
46
47     // ... now process those samples ...
```



```

48     processLastBrakeEvents(lastNActiveSamples);
49 }
50
51 int main(int argc, char** argv) {
52
53     /* Instance Specifier from model */
54     ara::core::InstanceSpecifier instspec {...}
55
56     auto handles = proxy::RadarServiceProxy::FindService(instspec);
57
58     if (!handles.empty()) {
59         /* we have at least one valid handle - we are not very particular
60          * here and take the first one to create our proxy */
61         myRadarProxy = std::make_unique<proxy::RadarServiceProxy>(handles
62 [0]);
63
64         /* we are interested in receiving the event "BrakeEvent" - so we
65          * subscribe for it. We want to access up to 10 events, since our
66          * sample algo averages over at most 10.*/
67         myRadarProxy->BrakeEvent.Subscribe(10);
68
69         /* whenever new BrakeEvent events come in, we want be called, so we
70          * register a callback for it!
71          * Note: If the entity we would subscribe to, would be a field
72          * instead of an event, it would be crucial, to register our
73          * reception handler BEFORE subscribing, to avoid race conditions.
74          * After a field subscription, you would get instantly so called
75          * "initial events" and to be sure not to miss them, you should
76          care
77          * for that your reception handler is registered before.*/
78         myRadarProxy->BrakeEvent.SetReceiveHandler(
79 handleBrakeEventReception);
80     }
81
82     // ... wait for application shutdown trigger by application exec mgmt.
83 }

```

Listing 5.5: Sample Code how to access Events

5.3.5.6 Buffering Strategies

The following figure sketches a simple deployment, where we have a service providing an event, for which two different local adaptive SWCs have subscribed through their respective `ara::com` proxies/event wrappers.

As you can see in the picture both proxies have a local event cache. This is the cache, which gets filled via `GetNewSamples()`. What this picture also depicts is, that the service implementation sends its event data to a Communication Management buffer, which is apparently outside the process space of the service implementation — the picture here assumes, that this buffer is owned by kernel or it is realized as a shared memory between communicating proxies and skeleton or owned by a separate binding implementation specific “demon” process.

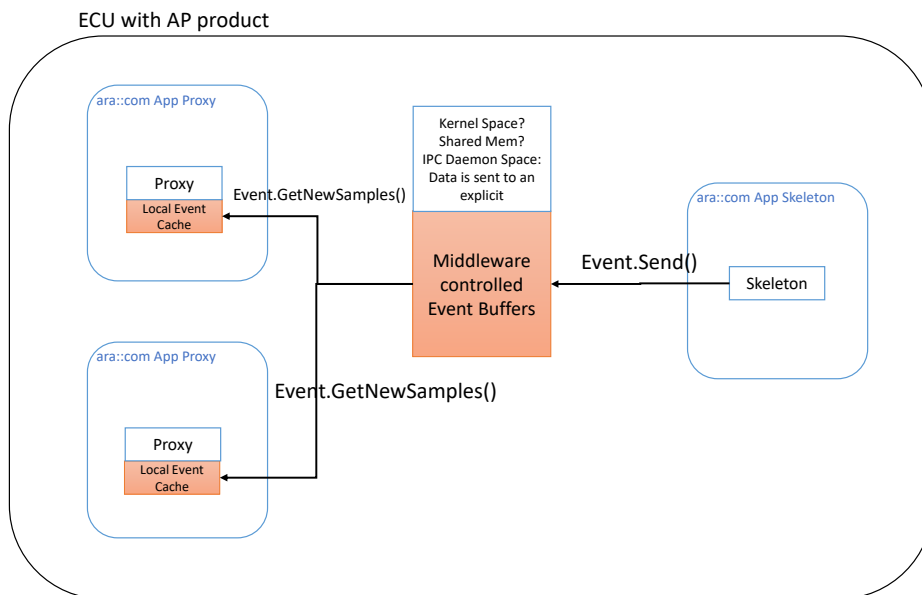


Figure 5.2: Event Buffering Approaches

The background of those assumptions made in the figure is the following: Adaptive applications are realized as processes with separated/protected memory/address spaces.

Event Data sent out by the service implementation (via the skeleton) cannot be buffered inside the service/skeleton process private address space: If that would be the case, event data access by the proxies would typically lead to context switches to the service application process.

Something, which we want to have total control over on service side via the [Method-CallProcessingMode](#) (see [Section 5.4.5](#)) and should therefore not be triggered by the communication behavior of arbitrary service consumers. Now let's have a rough look at the three different places, where the buffer, which is target for the "send event" might be located:

- **Kernel Space:** Data is sent to a memory region not mapped directly to an application process. This is typically the case, when binding implementation uses [IPC](#) primitives like pipes or sockets, where data written to such a primitive ends up in kernel buffer space.
- **Shared Memory:** Data is sent to a memory region, which is also directly readable from receivers/proxies. Writing/reading between different parties is synchronized specifically (lightweight with mem barriers or with explicit mutexes).
- **IPC-Daemon Space:** Data is sent to an explicit non-application process, which acts as a kind of demon for the IPC/binding implementation. Note, that technically this approach might be built on an [IPC](#) primitive like communication via kernel space or shared memory to get the data from service process to demon process.

Each of those approaches might have different pros and cons regarding flexibility/size of buffer space, efficiency in terms of access speed/overhead and protection against malicious access/writing of buffers. Therefore consideration of different constraints in an AP product and its use might lead to different solutions.

What shall be emphasized here in this example, is, that the AP product vendor is explicitly encouraged to use a reference based approach to access event data: The `ara::com` API of event wrapper intentionally models the access via `SamplePtr`, which are passed to the callbacks and not the value!

In those rather typical scenarios of 1:N event communication, this would allow to have inside the “Local Event Cache” not the event data values itself but pointers/references to the data contained in a central Communication Management buffer. Updating the local cache via `GetNewSamples()` could then be implemented not as a value copy but as reference updates.

To be honest: This is obviously a coarse grained picture of optimization possibilities regarding buffer usage! As hinted here ([Section 7.1](#)) data transferred to application processes must typically be de-serialized latest before first application access.

Since de-serialization has to be specific to the alignment of the consuming application the central sharing of an already de-serialized representation might be tricky. But at least you get the point, that the API design for event data access on the proxy/service consumer side gives room to apply event data sharing among consumers.

5.3.6 Methods

For each method the remote service provides, the proxy class contains a member of a method specific wrapper class.

In our example, we have three methods and the corresponding members have the name `Calibrate` (of type `methods::Calibrate`), `Adjust` (of type `methods::Adjust`) and `LogCurrentState` (of type `methods::LogCurrentState`). Just like the event classes the needed method classes of the proxy class are generated inside a specific namespace `methods`, which is contained inside the `proxy` namespace.

The method member in the proxy is used to call a method provided by the possibly remote service instance our proxy is connected to.

Let's have a look at the generated method class for our example — we pick out the `Adjust` method here:

```

1 class Adjust {
2     public:
3         /**
4          * For all output and non-void return parameters
5          * an enclosing struct is generated, which contains
6          * non-void return value and/or out parameters.
7          */
8         struct Output {
9             bool success;
10            Position effective_position;
11        };
12
13        /**
14         * \brief Operation will call the method.
15         *
16         * Using the operator the call will be made by the Communication
17         * Management and a future returned, which allows the caller to
18         * get access to the method result.
19         *
20         * \param[in] target_position See service description.
21         *
22         * \return A future containing Output struct
23         */
24        ara::core::Future<Output> operator()(const Position &target_position);
25    };

```

Listing 5.6: Proxy side Adjust Method Class

So the method wrapper class is not that complex. It just consists of two parts: An inner structure definition, which aggregates all OUT-/INOUT-parameters of the method, and a bracket operator, which is used to call the service method.

The operator contains all of the service methods IN-/INOUT-parameters as IN-parameters. That means INOUT-parameters in the abstract service method description are split in a pair of IN and OUT parameters in the `ara::com` API.

The return value of a call to a service method, which is **not** a “one-way method” is an `ara::core::Future`, where the template parameter is of the type of the inner `struct`, which aggregates all OUT-parameters of the method. More about this `ara::core::Future` in the following subsection.

5.3.6.1 One-Way aka Fire-and-Forget Methods

Before proceeding with the functionalities provided for “normal” methods, we briefly introduce “one-way methods” here as we already referred to this term in the previous section. `ara::com` supports a special flavor of a method, which we call “one-way” or “fire-and-forget”. Technically this is a method with only IN-params — no OUT-params and no raising of errors allowed. There is also no hand-shaking/synchronisation possible with the server! The client/caller therefore gets no feedback at all, whether the server/callee has processed a “one-way” call or not (except for a possible recoverable `Error` (`ComErrc::kNetworkBindingFailure`) in the returned `ara::core::Result` object).

There are communication patterns, where such a best-effort approach is fully sufficient. In this case such a “one-way/fire-and-forget” semantics is very light-weight from a resource perspective. If we look at the signature of such a method, we see, that it is simpler, than that from a regular method:

```
1 class LogCurrentState {
2     public:
3         /**
4          * \brief Operation will call the method.
5          *
6          * Using the operator the call will be made by the Communication
7          * Management.
8          *
9          * It is a one-way method, so no feedback (return value/out-parameter)
10         * is given, except for a possible error code kNetworkBindingFailure
11         * in the returned ara::core::Result object if a recoverable error is
12         * reported from the network binding.
13         */
14         ara::core::Result<void> operator() ();
15     };
```

Listing 5.7: Proxy side LogCurrentState Method Class

5.3.6.2 Event-Driven vs Polling access to method results

Similar to the access to event data described in the previous section ([Section 5.3.5](#)), we provide API support for an event-driven and polling-based approach also for accessing the results of a service method call.

The magic of differentiation between both approaches lies in the returned `ara::core::Future`: `ara::core::Future` is basically an extended version of the C++11/C++14 `std::future` class; see [2] for details.

Like in the event data access, event-driven here means, that the caller of the method (the application with the proxy instance) gets notified by the Communication Management implementation as soon as the method call result has arrived.

For a Communication Management implementation of `ara::com` this means, it has to setup some kind of waiting mechanism (WaitEvent) behind the scene, which gets woken up as soon as the method result becomes available, to notify the `ara::com` user. So how do the different usage patterns of the `ara::core::Future` work then?

Let's have a deeper look at our `ara::core::Future` and the interfaces it provides:

```

1 enum class future_status : uint8_t
2 {
3     ready,    ///< the shared state is ready
4     timeout, ///< the shared state did not become ready before the specified
               timeout has passed
5 };
6
7 template <typename T, typename E = ErrorCode>
8 class Future {
9     public:
10
11     Future() noexcept = default;
12     ~Future();
13
14     Future(const Future&) = delete;
15     Future& operator=(const Future&) = delete;
16
17     Future(Future&& other) noexcept;
18     Future& operator=(Future&& other) noexcept;
19
20     /**
21     * @brief Get the value.
22     *
23     * This function shall behave the same as the corresponding std::future
24     * function.
25     *
26     * @returns value of type T
27     * @error Domain:error the error that has been put into the
28     * corresponding Promise via Promise::SetError
29     *
30     */
31     T get();
32
33     /**

```

```

34     * @brief Get the result.
35     *
36     * Similar to get(), this call blocks until the value or an error is
37     * available. However, this call will never throw an exception.
38     *
39     * @returns a Result with either a value or an error
40     * @error Domain:error the error that has been put into the
41     * corresponding Promise via Promise::SetError
42     *
43     */
44     Result<T, E> GetResult() noexcept;
45
46     /**
47     * @brief Checks if the Future is valid, i.e. if it has a shared state.
48     *
49     * This function shall behave the same as the corresponding std::future
50     * function.
51     *
52     * @returns true if the Future is usable, false otherwise
53     */
54     bool valid() const noexcept;
55
56     /**
57     * @brief Wait for a value or an error to be available.
58     *
59     * This function shall behave the same as the corresponding std::future
60     * function.
61     */
62     void wait() const;
63
64     /**
65     * @brief Wait for the given period, or until a value or an error is
66     * available.
67     *
68     * This function shall behave the same as the corresponding std::future
69     * function.
70     *
71     * @param timeoutDuration maximal duration to wait for
72     * @returns status that indicates whether the timeout hit or if a value
73     * is available
74     */
75     template <typename Rep, typename Period>
76     future_status wait_for(std::chrono::duration<Rep, Period> const&
77     timeoutDuration) const;
78
79     /**
80     * @brief Wait until the given time, or until a value or an error is
81     * available.
82     *
83     * This function shall behave the same as the corresponding std::future
84     * function.
85     *
86     * @param deadline latest point in time to wait
87     * @returns status that indicates whether the time was reached or if a
88     * value is available
89     */

```



```

89     template <typename Clock, typename Duration>
90     future_status wait_until(const std::chrono::time_point<Clock, Duration
    >& deadline) const;
91
92     /**
93     * @brief Register a callable that gets called when the Future becomes
94     * ready.
95     *
96     * When @a func is called, it is guaranteed that get() and GetResult()
97     * will not block.
98     *
99     * @a func may be called in the context of this call or in the context
100    * of Promise::set_value() or Promise::SetError() or somewhere else.
101    *
102    * The return type of @a then depends on the return type of @a func
103    * (aka continuation).
104    *
105    * Let U be the return type of the continuation (i.e. a type equivalent
106    * to std::result_of_t<std::decay_t<F>(Future<T,E>>>).
107    * - If U is Future<T2,E2> for some types T2, E2, then the return
108    *   type of @a then() is Future<T2,E2>. This is known as implicit
109    *   Future unwrapping.
110    * - If U is Result<T2,E2> for some types T2, E2, then the return
111    *   type of @a then() is Future<T2,E2>. This is known as implicit
112    *   Result unwrapping.
113    * - Otherwise it is Future<U,E>.
114    *
115    * @param func a callable to register
116    * @returns a new Future instance for the result of the continuation
117    */
118    template <typename F>
119    auto then(F&& func) -> Future<SEE_COMMENT_ABOVE>;
120
121    /**
122    * @brief Return whether the asynchronous operation has finished.
123    *
124    * If this function returns true, get(), GetResult() and the wait calls
125    * are guaranteed not to block.
126    *
127    * The behavior of this function is undefined if valid() returns false.
128    *
129    * @returns true if the Future contains a value or an error, false
130    * otherwise
131    */
132    bool is_ready() const;
133 };

```

Listing 5.8: ara::core::Future Class

`ara::core::GetResult()` returns `Result` or an `Error` inside an `ara::core::Future` object and throws no exception. `get()` returns the corresponding `ara::core::Future` and / or throws exception.

See [2] chapter "Error handling" for detailed documentation of the error handling approaches in the Addaptive Platform.

Below is the sample of using "exception-based" approach to synchronously call a method:

```
1 using namespace ara::com;
2
3 int main() {
4     // some code to acquire a handle
5     // ...
6     RadarServiceProxy service(handle);
7     Future<Calibrate::Output> callFuture = service.Calibrate(
8         myConfigString);
9
10    /**
11     * Now we do a blocking get(), which will return in case the result
12     * (valid or exception) is received.
13     *
14     * If Calibrate could throw an exception and the service has set one,
15     * it would be thrown by get()
16     */
17    Calibrate::Output callOutput = callFuture.get();
18
19    // process callOutput ...
20    return 0;
21 }
```

Listing 5.9: Synchronous method call sample

In a nutshell: A synchronous call (from the viewpoint of the application developer) to a service method, simply consists of the `()`-operator call-syntax with a subsequent blocking `get()` call on the returned future.

There are other ways for the user to get a notification from the Communication Management implementation as soon as the method result is available beside resuming execution from a blocking call to `get()`:

- The variants of “wait”, which the `ara::core::Future` has taken over from `std::future`. They basically provide the functionality of a blocking wait for the fulfillment of the future.
- Registering a callback method via `then()`. This is one of the extensions to the `std::future`; see [2] for details.

The plain parameterless `wait()` variant has the same blocking semantics like `get()` — i.e. blocks till the future has a valid result (value or exception).

The variants of “wait”, where you either give a duration (`wait_for()`) or a target point in time (`wait_until()`) will return either if the future has a valid result or in case the timeout/deadline restriction has been met — therefore they both return `future_status` to allow distinction between those cases.

The last possibility to get notification of the result of the future (valid or exception) is by registering a callback method via `then()`. This is one of the extensions to the `ara::core::Future` over `std::future`.

As you can see, all the possibilities to get access to the future's method result we have discussed (and partly showed in examples) up to now — blocking “get”, all “wait” variants and “then” — are **event-driven**. I.e. the event of the arrival of the method result (or an error) leads to either resuming of a blocked user thread or call to a user provided function!

There are of course cases, where the `ara::com` users does not want his application (process) getting activated by some method-call return event at all! Think for a typical RT (real time) application, which must be in total control of its execution. We discussed this RT/polling use case already in the context of event data access already ([Section 5.3.5.3](#)). For method calls the same approach applies!

So we did foresee the following usage pattern with regards to `ara::core::Future`: After you have called the service method via the `()`-operator, you just use `ara::core::Future::is_ready()` to poll, whether the method call has been finished. This call is defined to be **non-blocking**. Sure, it might involve some syscall/context-switch (for instance to look into some kernel buffers), which is not for free, but it does not block!

After `ara::core::Future::is_ready()` has returned `true`, it is guaranteed that the next call to `ara::core::Future::get()` will NOT block, but immediately return either the valid value or throw an exception in case of error.

5.3.6.3 Canceling Method Result

There may be cases, where you already have called a service method via the `()`-operator, which returned you an `ara::core::Future`, but you are not interested in the result anymore.

It could even be the case, that you already have registered a callback via `ara::core::Future::then()` for it. Instead of just let things go and “ignore” the callback, you should tell the Communication Management explicitly.

This might free resources and avoid unnecessary processing load on the binding implementation level. Telling that you are not interested in the method call result anymore is simply done by letting the `ara::core::Future` go out of scope, so that its destructor gets called.

Call of the `dtor` of the `ara::core::Future` is a signal to the binding implementation, that any registered callback for this future shall not be called anymore, reserved/allocated memory for the method call result might be freed and event waiting mechanisms for the method result shall be stopped.

To trigger the call to the `dtor` you could obviously let the future go out of scope. Depending on the application architecture this might not be feasible, as you already

might have assigned the returned `ara::core::Future` to some variable with greater scope.

To solve this, the `ara::core::Future` is default-constructible. Therefore you simply overwrite the returned `ara::core::Future` in the variable with a default constructed instance as is shown in the example below:

```
1 using namespace ara::com;
2
3 Future<Calibrate::Output> calibrateFuture;
4
5 int main() {
6     // some code to acquire handle
7     // ...
8     RadarServiceProxy service(handle);
9     calibrateFuture = service.Calibrate(myConfigString);
10
11     /** ....
12      * Some state changes happened, which render the calibrate method
13      * result superfluous ...
14      *
15      * We force deletion by resetting our variable to a new default
16      * constructed Future.
17      */
18     calibrateFuture = Future<Calibrate::Output>();
19
20     // go on doing something ...
21     return 0;
22 }
```

Listing 5.10: Example of discarding a future

5.3.7 Fields

Conceptually a field has — unlike an event — a certain value at any time. That results in the following additions compared to an event:

- if a subscription to a field has been done, “immediately” current values are sent back to the subscriber in an event-like notification pattern.
- the current field value can be queried via a call to a `Get ()` method or could be updated via a `Set ()` method.

Note, that all the features a field provides are optionally: In the configuration (IDL) of your field, you decide, whether it has “on-change-notification”, `Get ()` or `Set ()`. In our example field (see below), we have all three mechanisms configured.

For each field the remote service provides, the proxy class contains a member of a field specific wrapper class. In our example the member has the name `UpdateRate` (of type `fields::UpdateRate`).

Just like the event and method classes the needed field classes of the proxy class are generated inside a specific namespace `fields`, which is contained inside the `proxy` namespace.

The explanation of fields has been intentionally put after the explanation of events and methods, since the field concept is roughly an aggregation of an event with correlated `get()/set()` methods. Therefore technically we also implement the `ara::com` field representation as a combination of `ara::com` event and method.

Consequently the field member in the proxy is used to

- call `Get ()` or `Set ()` methods of the field with exactly the same mechanism as regular methods
- access field update notifications in the form of events/event data, which are sent by the service instance our proxy is connected to with exactly the same mechanism as regular events

Let's have a look at the generated field class for our example UpdateRate field here:

```

1 class UpdateRate {
2     /**
3      * \brief Shortcut for the events data type.
4      */
5     using FieldType = uint32_t;
6
7     /**
8      * \brief See Events for details, as a field contains the possibility
9      * for notifications the details of the interfaces described there.
10    */
11    ara::core::Result<void> Subscribe(size_t maxSampleCount);
12    size_t GetFreeSampleCount() const noexcept;
13    ara::com::SubscriptionState GetSubscriptionState() const;
14    void Unsubscribe();
15    ara::core::Result<void> SetReceiveHandler(ara::com::EventReceiveHandler
16    handler);
17    ara::core::Result<void> UnsetReceiveHandler();
18    ara::core::Result<void> SetSubscriptionStateChangeHandler(ara::com::
19    SubscriptionStateChangeHandler handler);
20    void UnsetSubscriptionStateChangeHandler();
21    template <typename F>
22    ara::core::Result<size_t> GetNewSamples(
23        F&& f,
24        size_t maxNumberOfSamples = std::numeric_limits<size_t>::max());
25
26    /**
27     * The getter allows to request the actual value of the service
28     * provider.
29     *
30     * For a description of the future, see the method.
31     * It should behave like a Method.
32     */
33    ara::core::Future<FieldType> Get();
34
35    /**
36     * The setter allows to request the setting of a new value.
37     * It is up to the Service Provider to accept the request or modify it.
38     * The new value shall be sent back to the requester as response.
39     *
40     * For a description of the future, see the method.
41     * It should behave like a Method.
42     */
43    ara::core::Future<FieldType> Set(const FieldType& value);
44 };

```

Listing 5.11: Proxy side UpdateRate Field Class

There is nothing more to be described here. For documentation of the mechanisms of event-like part of the field have a look at [Section 5.3.5](#) and for documentation of the method-like part of the field have a look at [Section 5.3.6](#).

5.3.8 Triggers

`Triggers` are simply a class of dataless events. Which means that all the documentation in [Section 5.3.5](#) is also applicable for `Triggers`, except for the data related part, which is described in this section.

From `Proxy` point of view, subscribing to a `Trigger` is the same as described in [Section 5.3.5.1](#) without the local cache part.

```

1     /**
2     * \brief The application expects the CM to subscribe the event.
3     */
4     ara::core::Result<void> Subscribe();

```

A `Subscribe()` call will inform Communication Management for receiving `Trigger` updates.

The other subscription capabilities (e.g. `SubscribeChangeSetSubscriptionHandler()`, `Unsubscribe()`) are the same as for `Events`.

A big difference from `Events` is receiving `Trigger` updates.

In contrary to `Events`, where we were interested in the received data, for `Triggers` we are only interested in the number of `Triggers` received since last check. Therefore `GetNewTriggers()` is more simple than `GetNewSamples()`:

```

1     /**
2     * \brief Get number of triggeres if any recieved.
3     *
4     * \return Result, which indicates the number of new triggers
5     * (zero means no new trigger has been recieved).
6     */
7     size_t GetNewTriggers();

```

Where the return value has the number of received `Triggers` that took place since the last call to `GetNewTriggers()`.

Polling-based and Event-driven access, as mentioned in [Section 5.3.5.5](#), is also supported for `Triggers`

5.4 Skeleton Class

The Skeleton class is generated from the [SI](#) description of the AUTOSAR meta model. `ara::com` does standardize the interface of the generated Skeleton class. The toolchain of an AP product vendor will generate a Skeleton implementation class exactly implementing this interface.

The generated Skeleton class is an abstract class. It cannot be instantiated directly, because it does not contain implementations of the service methods, which the service shall provide. Therefore the service implementer has to subclass the skeleton and provide the service method implementation within the subclass.

Note: Equal to the Proxy class the interfaces the Skeleton class has to provide are defined by `ara::com`, a generic (product independent) generator could generate an abstract class or a mock class against which the application developer could implement his service provider application. This perfectly suits the platform vendor independent development of Adaptive AUTOSAR SWCs.

`ara::com` expects skeleton related artifacts inside a namespace "skeleton". This namespace is typically included in a namespace hierarchy deduced from the service definition and its context.

5.4.1 Skeleton Class API's

- `OfferService()`
- `StopOfferService()`
- `Send()`
- `Allocate()`
- `ProcessNextMethodCall()`
- `RegisterGetHandler()`
- `RegisterSetHandler()`
- `Field::Update()`

5.4.2 RadarService Skeleton Class Example

```

1 class RadarServiceSkeleton {
2     public:
3         /**
4          * Ctor taking instance identifier as parameter and having default
5          * request processing mode kEvent.
6          */
7         RadarServiceSkeleton(ara::com::InstanceIdIdentifier instanceId,
8                             ara::com::MethodCallProcessingMode mode =
9                             ara::com::MethodCallProcessingMode::kEvent);
10
11        /**
12         * Exception-less ctor taking instance identifier as parameter
13         * and having default request processing mode kEvent.
14         */
15        static ara::core::Result<RadarServiceSkeleton> Create(
16            const ara::core::InstanceIdIdentifier &instanceID,
17            ara::com::MethodCallProcessingMode mode =
18            ara::com::MethodCallProcessingMode::kEvent) noexcept;
19
20        /**
21         * Ctor taking instance identifier container as parameter and having
    
```



```

22     * default request processing mode kEvent.
23     * This specifically supports multi-binding.
24     */
25 RadarServiceSkeleton(ara::com::InstanceIdentifierContainer instanceIds,
26 ara::com::MethodCallProcessingMode mode =
27 ara::com::MethodCallProcessingMode::kEvent);
28
29 /**
30     * Exception-less ctor taking instance identifier container as
parameter
31     * and having default request processing mode kEvent.
32     */
33 static ara::core::Result<RadarServiceSkeleton> Create(
34 const ara::core::InstanceIdentifierContainer &instanceIDs,
35 ara::com::MethodCallProcessingMode mode =
36 ara::com::MethodCallProcessingMode::kEvent) noexcept;
37
38 /**
39     * Ctor taking instance specifier as parameter and having default
40     * request processing mode kEvent.
41     */
42 RadarServiceSkeleton(ara::core::InstanceSpecifier instanceSpec,
43 ara::com::MethodCallProcessingMode mode =
44 ara::com::MethodCallProcessingMode::kEvent);
45
46 /**
47     * Exception-less ctor taking instance specifier as parameter and
having default
48     * request processing mode kEvent.
49     */
50 static ara::core::Result<RadarServiceSkeleton> Create(
51 const ara::core::InstanceSpecifier &instanceSpec,
52 ara::com::MethodCallProcessingMode mode =
53 ara::com::MethodCallProcessingMode::kEvent) noexcept;
54
55 /**
56     * skeleton instances are nor copy constructible.
57     */
58 RadarServiceSkeleton(const RadarServiceSkeleton& other) = delete;
59
60 /**
61     * skeleton instances are nor copy assignable.
62     */
63 RadarServiceSkeleton& operator=(const RadarServiceSkeleton& other) =
delete;
64
65 /**
66     * The Communication Management implementer should care in his dtor
67     * implementation, that the functionality of StopOfferService()
68     * is internally triggered in case this service instance has
69     * been offered before. This is a convenient cleanup functionality.
70     */
71 ~RadarServiceSkeleton();
72
73 /**
74     * Offer the service instance.

```

```

75     * method is idempotent - could be called repeatedly.
76     */
77     ara::core::Result<void> OfferService();
78
79     /**
80     * Stop Offering the service instance.
81     * method is idempotent - could be called repeatedly.
82     *
83     * If service instance gets destroyed - it is expected that the
84     * Communication Management implementation calls StopOfferService()
85     * internally.
86     */
87     void StopOfferService();
88
89     /**
90     * For all output and non-void return parameters
91     * an enclosing struct is generated, which contains
92     * non-void return value and/or out parameters.
93     */
94     struct CalibrateOutput {
95         bool result;
96     };
97
98     /**
99     * For all output and non-void return parameters
100    * an enclosing struct is generated, which contains
101    * non-void return value and/or out parameters.
102    */
103    struct AdjustOutput {
104        bool success;
105        Position effective_position;
106    };
107
108    /**
109    * This fetches the next call from the Communication Management
110    * and executes it. The return value is a bool which is set to true
111    * if further methods calls are pending to be processed.
112    * Only available in polling mode.
113    */
114    bool ProcessNextMethodCall();
115
116    /**
117    * \brief Public member for the BrakeEvent
118    */
119    events::BrakeEvent BrakeEvent;
120
121    /**
122    * \brief Public member for the UpdateRate
123    */
124    fields::UpdateRate UpdateRate;
125
126    /**
127    * The following methods are pure virtual and have to be implemented
128    */
129    virtual ara::core::Future<CalibrateOutput> Calibrate(
130    std::string configuration) = 0;

```

```
131     virtual ara::core::Future<AdjustOutput> Adjust(  
132     const Position& position) = 0;  
133     virtual void LogCurrentState() = 0;  
134 };
```

Listing 5.12: RadarService Skeleton

5.4.3 Instantiation (Constructors)

As you see in the example code of the `RadarServiceSkeleton` above, the skeleton class, from which the service implementer has to subclass his service implementation, provides three different `ctor` variants, which basically differ in the way, how the instance identifier to be used is determined.

Since you could deploy many different instances of the same type (and therefore same skeleton class) it is straightforward, that you have to give an instance identifier upon creation. This identifier has to be unique. In the exception-less creation of a service Skeleton using the named constructor approach, a static member function, `Create()`, checks if the provided identifier is not unique, or other errors. If an error is discovered, an error code is set in the returned `ara::core::Result`. Otherwise the created Skeleton instance is returned.

If a new instance shall be created with the same identifier, the existing instance needs to be destroyed before.

Exactly for this reason the skeleton class (just like the proxy class) does neither support copy construction nor copy assignment! Otherwise two "identical" instances would exist for some time with the same instance identifier and routing of method calls would be non-deterministic.

The different variants of `ctors` regarding instance identifier definition reflect their different natures, which are described in [Section 4.8.1](#).

- variant with `ara::com::InstanceIdentifier`: Service instance will be created with exactly one binding specific instance identifier.
- variant with `ara::com::InstanceIdentifierContainer`: Service instance will be created with bindings to multiple distinct instance identifiers. This is mentioned as "multi-binding" throughout this document and also explained in more detail in [Section 7.3](#)
- variant with `ara::core::InstanceSpecifier`: Service instance will be created with bindings to the instance identifier(s) found after "service manifest" lookup with the given `ara::core::InstanceSpecifier`. Note, that this could also imply a "multi-binding" as the integrator could have mapped the given `ara::core::InstanceSpecifier` to multiple technical/binding specific instance identifiers within the "service manifest".

The second parameter of the `ctors` of type `ara::com::MethodCallProcessingMode` has a default value and is explained in detail in [Section 5.4.5](#).

Note: Directly after creation of an instance of the subclass implementing the skeleton, this instance will not be visible to potential consumers and therefore no method will be called on it. This is only possible after the service instance has been made visible with the `OfferService` API (see below).

5.4.4 Offering Service instance

The skeleton provides the method `OfferService()`. After you — as application developer for the service provider side — have instantiated your custom service implementation class and initialized/set up your instance to a state, where it is now able to serve requests (method calls) and provide events to subscribing consumers, you will call this `OfferService()` method on your instance.

From this point in time, where you call it, method calls might be dispatched to your service instance — even if the call to `OfferService()` has not yet returned.

If you decide at a certain point (maybe due to some state changes), that you do not want to provide the service anymore, you call `StopOfferService()` on your instance. The contract here is: After `StopOfferService()` has returned no further method calls will be dispatched to your service instance.

For sanity reasons `ara::com` has the requirement for the AP vendors implementation of the skeleton `dtor`, that it internally does a `StopOfferService()` too, if the instance is currently offered.

So — “stop offer” needs only be called on an instance which lives on, and during its lifetime it switches between states where it is visible and provides its service, and states where it does not provide the service.

```

1 using namespace ara::com;
2
3 /**
4  * Our implementation of the RadarService -
5  * subclass of RadarServiceSkeleton
6  */
7 class RadarServiceImpl;
8
9 int main(int argc, char** argv) {
10     // read instanceId from commandline
11     ara::core::string_view instanceIdStr(argv[1]);
12     RadarServiceImpl myRadarService(InstanceIdentifier(instanceIdStr));
13
14     // do some service specific initialization here ....
15     myRadarService.init();
16
17     // now service instance is ready -> make it visible/available
18     myRadarService.OfferService();
19
20     // go into some wait state in main thread - waiting for AppExecMgr
21     // signals or the like ....
22

```

```

23     return 0;
24 }

```

Listing 5.13: Example of RadarService Init and Offer

5.4.5 Polling and event-driven processing modes

Now let's come to the point, where we deliver on the promise to support event-driven and polling behavior also on the service providing side. From the viewpoint of the service providing instance — here our skeleton/skeleton subclass instance — requests (service method or field getter/setter calls) from service consumers may come in at arbitrary points in time.

In a purely event-driven setup, this would mean, that the Communication Management generates corresponding call events and transforms those events to concrete method calls to the service methods provided by the service implementation.

The consequences of this setup are clear:

- general reaction to a service method call might be fast, since the latency is only restricted by general machine load and intrinsic IPC mechanism latency.
- rate of context switches to the OS process containing the service instance might be high and non-deterministic, decreasing overall throughput.

As you see — there are pros and cons for an event-driven processing mode at the service provider side. However, we do support such a processing mode with `ara::com`. The other bookend we do support, is a pure polling style approach. Here the application developer on the service provider side explicitly calls an `ara::com` provided API to process explicitly **one** call event.

With this approach we again support the typical RT-application developer. His application gets typically activated due to a low jitter cyclical alarm.

When his application is active, it checks event queues in a non-blocking manner and decides explicitly how many of those accumulated (since last activation time) events it is willing to process. Again: Context switches/activations of the application process are only accepted by specific (RT) timers. Asynchronous communication events shall **not** lead to an application process activation.

So how does `ara::com` allow the application developer to differentiate between those processing modes? The behavior of a skeleton instance is controlled by the second parameter of its `ctor`, which is of type `ara::com::MethodCallProcessingMode`.

```

1 /**
2  * Request processing modes for the service implementation side
3  * (skeleton).
4  *
5  * \note Should be provided by platform vendor exactly like this.
6  */
7 enum class MethodCallProcessingMode { kPoll, kEvent, kEventSingleThread };

```

That means the processing mode is set for the entire service instance (i.e. all its provided methods are affected) and is fix for the whole lifetime of the skeleton instance. The default value in the `ctor` is set to `kEvent`, which is explained below.

5.4.5.1 Polling Mode

If you set it to `kPoll`, the Communication Management implementation will not call any of the provided service methods asynchronously!

If you want to process the next (assume that there is a queue behind the scenes, where incoming service method calls are stored) pending service-call, you have to call the following method on your service instance:

```
1 /**
2  * This fetches the next call from the Communication Management
3  * and executes it.
4  * Only available in polling mode.
5  */
6 bool ProcessNextMethodCall();
```

We are using the mechanism of a `bool` return value to immediately return a result. A simple use case for a typical RT application could be:

- RT application gets scheduled.
- It calls `ProcessNextMethodCall` and obtains its result immediately after, without any context switches, the first outstanding request processing has finished.
- Finally, the RT application decides whether there is enough time left for serving a subsequent service method. If so, it calls another `ProcessNextMethodCall`.

Sure - this simple example assumes, that the RT application knows worst case runtime of its service methods (and its overall time slice), but this is not that unlikely!

The returned `bool` value is set to `true` by the Communication Management in case there really was an outstanding request in the queue, which has been dispatched, otherwise it is set to `false`.

This is a somewhat comfortable indicator to the application developer, not to call repeatedly `ProcessNextMethodCall` although the request queue is empty. So calling `ProcessNextMethodCall` directly after a previous call returned `false` might most likely do nothing (except that incidentally in this minimal time frame a new request came in).

Please note polling mode has implications on AP products based on typical operating systems. Ruling out context switches to a process (containing a service implementation) caused by Communication Management events (incoming service method calls) means also: There are constraints for the location of the queue, which has to collect the service method call requests until they are consumed by the polling service implementation.

The queue must be realized either outside of the address space of the service provider application or it must be located in a shared memory like location, so that the sending part is able to write directly into the queue. In comparison to a shared memory solution the access from the polling service provider to below queue locations might come with higher costs/latency. Typical solutions of placing the queue outside of the service provider address space would be:

- Kernel space: If the binding implementation would use socket or pipe mechanisms, the kernel buffers being the target of the write-call would resemble the queue. Adapting/configuring maximal sizes of those buffers might in typical OS mean recompiling the kernel.
- User address space of a different binding/Communication Management demonstration application: Buffer space allocation for queues allocated within user space could typically be done more dynamic/flexible

5.4.5.2 Event-Driven Mode

If you set the processing mode to `kEvent` or `kEventSingleThread`, the Communication Management implementation will dispatch events asynchronously to the service method implementations at the time the service call from the service consumer comes in.

Opposed to the `kPoll` mode, here the service consumer implicitly controls/triggers service provider process activations with their method calls!

What is then the difference between `kEvent` and `kEventSingleThread`? `kEvent` means, that the Communication Management implementation may call the service method implementations concurrently.

That means for our example: If — at the same point in time — one call to method `Calibrate` and two calls to method `Adjust` arrive from different service consumers, the Communication Management implementation is allowed to take three threads from its internal thread-pool and do those three calls for the two service methods concurrently.

On the contrary the mode `kEventSingleThread` assures, that on the service instance only one service method at a time will be called by the Communication Management implementation.

That means, Communication Management implementation has to queue incoming service method call events for the same service instance and dispatch them one after the other.

Why did we provide those two variants? From a functional viewpoint only `kEvent` would have been enough! A service implementation, where certain service methods could not run concurrently, because of shared data/consistency needs, could simply do its synchronization (e.g. via `std::mutex`) on its own!

The reason is “efficiency”. If you have a service instance implementation, which has extensive synchronization needs, i.e. would synchronize almost all service method calls anyways, it would be a total waste of resources, if the Communication Management would “spend” N threads from its thread-pool resources, which directly after get a hard sync, sending N-1 of it to sleep.

For service implementations which lie in between — i.e. some methods can be called concurrently without any sync needs, some methods need at least partially synchronization — the service implementer has to decide, whether he uses `kEvent` and does synchronization on top on his own (possibly optimizing latency, responsiveness of his service instance) or whether he uses `kEventSingleThread`, which frees him from synchronizing on his own (possibly optimizing ECU overall throughput).

5.4.6 Methods

Service methods on the skeleton side are abstract methods, which have to be overwritten by the service implementation sub-classing the skeleton. Let's have a look at the `Adjust` method of our service example:

```
1 /**
2  * For all output and non-void return parameters
3  * an enclosing struct is generated, which contains
4  * non-void return value and/or out parameters.
5  */
6 struct AdjustOutput {
7     bool success;
8     Position effective_position;
9 };
10
11 virtual ara::core::Future<AdjustOutput> Adjust(
12     const Position& position) = 0;
```

Listing 5.14: Skeleton side Adjust method

The IN-parameters from the abstract definition of the service method are directly mapped to method parameters of the skeletons abstract method signature.

In this case it's the `position` argument from type `Position`, which is — as it is a non-primitive type — modeled as a “const ref”¹.

The interesting part of the method signature is the return type. The implementation of the service method has to return our extensively discussed `ara::core::Future`.

The idea is simple: We do not want to force the service method implementer to signal the finalization of the service method with the simple return of this “entry point” method!

Maybe the service implementer decides to dispatch the real processing of the service call to a central worker-thread pool! This would then be really ugly, when the “entry point” methods return would signal the completion of the service call to the Communication Management.

Then — in our worker thread pool scenario — we would have to block into some kind of wait point inside the service method and wait for some notification from the worker thread, that he has finished and only then we would return from the service method.

In this scenario we would have a blocked thread inside the service-method! From the viewpoint of efficient usage of modern multi-core CPUs this is not acceptable.

The returned `ara::core::Future` contains a structure as template parameter, which aggregates all the OUT-parameters of the service call.

The following two code examples show two variants of an implementation of `Adjust`. In the first variant the service method is directly processed synchronously in the method

¹The referenced object is provided by the Communication Management implementation until the service method call has set its promise (valid result or error). If the service implementer needs the referenced object beyond that, he has to make a copy.

body, so that an `ara::core::Future` with an already set result is returned, while in the second example, the work is dispatched to an asynchronous worker, so that the returned `ara::core::Future` may not have a set result at return.

```

1 using namespace ara::com;
2
3 /**
4  * Our implementation of RadarService
5  */
6 class RadarServiceImpl : public RadarServiceSkeleton {
7     public:
8
9         Future<AdjustOutput> Adjust(const Position& position)
10        {
11            ara::core::Promise<AdjustOutput> promise;
12
13            // calling synchronous internal adjust function, which delivers
14            results
15            struct AdjustOutput out = doAdjustInternal(
16                position,
17                &out.effective_position);
18            promise.set_value(out);
19
20            // we return a future from an already set promise...
21            return promise.get_future();
22        }
23     private:
24
25         AdjustOutput doAdjustInternal(const Position& position) {
26             // ... implementation
27         }
28 }
    
```

Listing 5.15: Example of returning Future with already set result

As you see in the example above: Inside the body of the service method an internal method is called, which does the work synchronously. I.e. after the return of “doAdjustInternal” in `out` the attributes, which resemble the service methods out-params are set. Then this `out` value is set at the `ara::core::Promise` and then the `Future` created from the `Promise` is returned.

This has the effect that the caller, who gets this `Future` as return, can immediately call `Future::get()`, which would not block, but immediately return the `AdjustOutput`.

Now let's have a look at the asynchronous worker thread variant:

```

1 using namespace ara::com;
2
3 /**
4  * Our implementation of the RadarService
5  */
6 class RadarServiceImpl : public RadarServiceSkeleton {
7
8     public:
9         Future<AdjustOutput> Adjust(const Position& position)
10        {
11            ara::core::Promise<AdjustOutput> promise;
12            auto future = promise.get_future();
13
14            // asynchronous call to internal adjust function in a new Thread
15            std::thread t(
16                [this] (const Position& pos, ara::core::Promise prom) {
17                    prom.set_value(doAdjustInternal(pos));
18                },
19                std::cref(position), std::move(promise)).detach();
20
21            // we return a future, which might be set or not at this point...
22            return future;
23        }
24
25     private:
26         AdjustOutput doAdjustInternal(const Position& position) {
27             // ... implementation
28         }
29 }

```

Listing 5.16: Example of returning Future with possibly unset result

In this example, “doAdjustInternal” is called within a different asynchronous thread. In this case we wrapped the call to “doAdjustInternal” inside a small lambda, which does the job of setting the value to the [Promise](#).

5.4.6.1 One-Way aka Fire-and-Forget Methods

“One-way/fire-and-forget” methods on the server/skeleton side do have (like on the proxy side) a simpler signature compared to normal methods. Since there is no feedback possible/needed towards the caller from the server it is a method that returns an [ara::core::Result<void>](#). The [Result](#) type is used to return an error code in case a recoverable local network binding failure occurs in the Communication Management when calling the method.

```

1     virtual ara::core::Result<void> LogCurrentState() = 0;

```

5.4.6.2 Raising Application Errors

Whenever on the implementation side of a service method, an `ApplicationError` — according to the interface description — is detected, the `ErrorCode` representing this `ApplicationError` simply has to be stored into the `Promise`, from which the `Future` is returned to the caller:

```
1 using namespace ara::com;
2 using namespace com::mycompany::division::radarservice;
3
4 /**
5
6     Our implementation of the RadarService
7     */
8     class RadarServiceImpl : public RadarServiceSkeleton {
9
10 public:
11 Future<CalibrateOutput> Calibrate(const std::string& configuration)
12 {
13     ara::core::Promise<CalibrateOutput> promise;
14     auto future = promise.get_future();
15
16     // we check the given configuration arg
17     if (!checkConfigString(configuration))
18     { // given arg is invalid:
19         // assume that in ARXMLs we have ErrorDomain with name SpecificErrors
20         // which contains InvalidConfigString error.
21         // Note that numeric error code will be casted to ara::core::ErrorCode
22         // implicitly.
23         promise.SetError(SpecificErrorsErrc::InvalidConfigString);
24     }
25
26     else
27     { ... }
28
29     // we return a future with a potentially set exception
30     return future;
31 }
32
33 private:
34 bool checkConfigString(const std::string& config);
35
36 std::string curValidConfig_;
37 }
```

Listing 5.17: Returning Future with possibly set exception

In this example, the implementation of “Calibrate” detects, that the given configuration string argument is invalid and sets the corresponding exception to the `Promise`.

5.4.7 Events

On the skeleton side the service implementation is in charge of notifying about occurrence of an event. As shown in [Listing 5.12](#) the skeleton provides a member of an event wrapper class per each provided event. The event wrapper class on the skeleton/event provider side looks obviously different than on the proxy/event consumer side.

On the service provider/skeleton side the service specific event wrapper classes are defined within the namespace `event` directly beneath the namespace `skeleton`. Let's have a deeper look at the event wrapper for our example event `BrakeEvent`:

```

1 class BrakeEvent {
2     public:
3         /**
4          * Shortcut for the events data type.
5          */
6         using SampleType = RadarObjects;
7
8         ara::core::Result<void> Send(const SampleType &data);
9
10        ara::core::Result<ara::com::SampleAllocateePtr<SampleType>> Allocate();
11
12        /**
13         * After sending data you loose ownership and can't access
14         * the data through the SampleAllocateePtr anymore.
15         * Implementation of SampleAllocateePtr will be with the
16         * semantics of std::unique_ptr (see types.h)
17         */
18        ara::core::Result<void> Send(ara::com::SampleAllocateePtr<SampleType>
19        data);
20    };

```

Listing 5.18: Skeleton side of BrakeEvent class

The using directive — analogue to the Proxy side — just introduces the common name `SampleType` for the concrete data type of the event. We provide two different variants of a “Send” method, which is used to send out new event data. The first one takes a reference to a `SampleType`.

This variant is straight forward: The event data has been allocated somewhere by the service application developer and is given via reference to the binding implementation of `Send()`.

After the call to send returns, the data might be removed/changed on the caller side. The binding implementation will make a copy in the call.

The second variant of ‘Send’ also has a parameter named “data”, but this is now of a different type `ara::com::SampleAllocateePtr<SampleType>`. According to our general approach to only provide abstract interfaces and eventually provide a proposed mapping to existing C++ types (see [Section 4.6](#)), this pointer type that we introduce here, shall behave like a `std::unique_ptr<T>`.

That roughly means: Only one party can hold the pointer - if the owner wants to give it away, he has to explicitly do it via `std::move`. So what does this mean here? Why do we want to have `std::unique_ptr<T>` semantics here?

To understand the concept, we have to look at the third method within the event wrapper class first:

```
1     ara::com::SampleAllocateePtr<SampleType> Allocate();
```

The event wrapper class provides us here with a method to allocate memory for one sample of event data. It returns a smart pointer `ara::com::SampleAllocateePtr<SampleType>`, which points to the allocated memory, where we then can write an event data sample to. And this returned smart pointer we can then give into an upcoming call to the second version of “Send”.

So — the obvious question would be — why should I let the binding implementation do the memory allocation for event data, which I want to notify/send to potential consumers? The answer simply is: Possibility for optimization of data copies.

The following “over-simplified” example makes things clearer: Let’s say the event, which we talk about here (of type `RadarObjects`), could be quite big, i.e. it contains a vector, which can grow very large (say hundreds of kilobytes). In the first variant of “Send”, you would allocate the memory for this event on your own on the heap of your own application process.

Then — during the call to the first variant of “Send” — the binding implementation has to copy this event data from the (private) process heap to a memory location, where it would be accessible for the consumer. If the event data to copy is very large and the frequency of such event occurrences is high, the sheer runtime of the data copying might hurt.

The idea of the combination of `Allocate()` and the second variant to send event data (`Send(SampleAllocateePtr<SampleType>)`) is to eventually avoid this copy!

A smart binding implementation might implement the `Allocate()` in a way, that it allocates memory at a location, where writer (service/event provider) and reader (service/event consumer) can both directly access it! So an `ara::com::SampleAllocateePtr<SampleType>` is a pointer, which points to memory nearby the receiver.

Such locations, where two parties can both have direct access to, are typically called “shared memory”. The access to such regions should — for the sake of data consistency — be synchronized between readers and writers.

This is the reason, that the `Allocate()` method returns such a smart pointer with the aspects of single/solely user of the data, which it points to: After the potential writer (service/event provider side) has called `Allocate()`, he can access/write the data pointed to as long as he hands it over to the second send variant, where he explicitly gives away ownership!

This is needed, because after the call, the readers will access the data and need a consistent view of it.

```

1 using namespace ara::com;
2
3 // our implementation of RadarService - subclass of RadarServiceSkeleton
4 RadarServiceImpl myRadarService;
5
6 /**
7  * Handler called at occurrence of a BrakeEvent
8  */
9 void BrakeEventHandler() {
10
11     // let the binding allocate memory for event data...
12     SampleAllocateePtr<BrakeEvent::SampleType> curSamplePtr =
13     myRadarService.BrakeEvent.Allocate();
14
15     // fill the event data ...
16     curSamplePtr->active = true;
17     fillVector(curSamplePtr->objects);
18
19     // Now notify event to consumers ...
20     myRadarService.BrakeEvent.Send(std::move(curSamplePtr));
21
22     // Now any access to data via curSamplePtr would fail -
23     // we've given up ownership!
24 }

```

Listing 5.19: Event Allocate/Send sample

5.4.8 Fields

On the skeleton side the service implementation is in charge of

- updating and notifying about changes of the value of a field.
- serving incoming `Get ()` calls.
- serving incoming `Set ()` calls.

As shown in [Listing 5.12](#) the skeleton provides a member of a field wrapper class per each provided field. The field wrapper class on the skeleton/field provider side looks obviously different than on the proxy/field consumer side.

On the service provider/skeleton side the service specific field wrapper classes are defined within the namespace `fields` directly beneath the namespace `skeleton`. Let's have a deeper look at the field wrapper in case of our example event `UpdateRate`:

```

1 class UpdateRate {
2     public:
3
4     using FieldType = uint32_t;
5
6     /**
7      * Update equals the send method of the event. This triggers the
8      * transmission of the notify (if configured) to
9      * the subscribed clients.
10     *
11     * In case of a configured Getter, this has to be called at least
12     * once to set the initial value.
13     */
14     ara::core::Result<void> Update(const FieldType& data);
15
16     /**
17     * Registering a GetHandler is optional. If registered the function
18     * is called whenever a get request is received.
19     *
20     * If no Getter is registered ara::com is responsible for responding
21     * to the request using the last value set by update.
22     *
23     * This implicitly requires at least one call to update after
24     * initialization of the Service, before the service
25     * is offered. This is up to the implementer of the service.
26     *
27     * The get handler shall return a future.
28     */
29     ara::core::Result<void> RegisterGetHandler(std::function<ara::core::
30     Future<FieldType>()> getHandler);
31
32     /**
33     * Registering a SetHandler is mandatory, if the field supports it.
34     * The handler gets the data the sender requested to be set.
35     * It has to validate the settings and perform an update of its
36     * internal data. The new value of the field should than be set
37     * in the future.
38     */

```

```
38     * The returned value is sent to the requester and is sent via
    notification to all subscribed entities.
39     */
40     ara::core::Result<void> RegisterSetHandler(std::function<ara::core::
    Future<FieldType>(const FieldType& data)> setHandler);
41 };
```

Listing 5.20: Skeleton side UpdateRate Class

The using directive — again as in the Event Class and on the Proxy side — just introduces the common name `FieldType` for the concrete data type of the field.

We provide an `Update` method by which the service implementer can update the current value of the field.

It is very similar to the simple/first variant of the `Send` method of the event class: The field data has been allocated somewhere by the service application developer and is given via reference to the binding implementation of `Update`. After the call to `Update` returns, the data might be removed/alterd on the caller side.

The binding implementation will make a (typically serialized) copy in the call.

In case “on-change-notification” is configured for the field, notifications to subscribers of this field will be triggered by the binding implementation in the course of the `Update` call.

5.4.8.1 Registering Getters

The `RegisterGetHandler` method provides the possibility to register a method implementation by the service implementer, which gets then called by the binding implementation on an incoming `Get()` call from any proxy instance.

The `RegisterGetHandler` method in the generated skeleton does **only** exist in case availability of “field getter” has been configured for the field in the IDL!

Registration of such a “GetHandler” is fully optional! Typically there is no need for a service implementer to provide such a handler. The binding implementation always has access to the latest value, which has been set via `Update`. So any incoming `Get()` call can be served by the Communication Management implementation standalone.

A theoretical reason for a service implementer to still provide a “GetHandler” could be: Calculating the new/current value of a field is costly/time consuming. Therefore the service implementer/field provider wants to defer this process until there is really need for that value (indicated by a getter call). In this case he could calculate the new field value within its “GetHandler” implementation and give it back via the known `ara::com` promise/future pattern.

If you look at the bigger picture, then such a setup with the discussed intention, where the service implementer provides and registers a “GetHandler” will not really make sense, if the field is configured with “on-change-notification”, too.

In this case, new subscribers will get potentially outdated field values on subscription, since updating of the field value is deferred to the explicit call of a “GetHandler”.

You also have to keep in mind: In such a setup, with enabled “on-change-notification” together with a registered “GetHandler” the Communication Management implementation will **not** automatically care for, that the value the developer returns from the “GetHandler” will be synchronized with value, which subscribers get via “on-change-notification” event!

If the implementation of “GetHandler” does not internally call `Update()` with the same value, which it will deliver back via `ara::com` promise, then the field value delivered via “on-change-notification” event will differ from the value returned to the `Get()` call. I.e. the Communication Management implementation will not automatically/internally call `Update()` with the value the “GetHandler” returned.

Bottom line: Using `RegisterGetHandler` is rather an exotic use case and developers should be aware of the intrinsic effect.

Additionally a user provided “GetHandler”, which only returns the current value, which has already been updated by the service implementation via `Update()`, is typically very inefficient! The Communication Management then has to call to user space and to additionally apply field serialization of the returned value at any incoming `Get()` call.

Both things could be totally “optimized away” if the developer does not register a “GetHandler” and leaves the handling of `Get()` calls entirely to the Communication Management implementation.

5.4.8.2 Registering Setters

Opposed to the `RegisterGetHandler` the `RegisterSetHandler` API has to be called by the service implementer in case it exists (i.e. field has been configured with setter support).

The reason, that we decided to make the registration of a “SetHandler” mandatory is simple: We expect, that the server implementation will always need to check the validity of a new/updated field values set by any anonymous client.

A look at the signature of the “SetHandler” `std::function<ara::core::Future<FieldType>(const FieldType& data)>` reveals that the registered handler does get the new value as input argument and is expected to return also a value. The semantic behind this is: In case the “SetHandler” always has to return the effective (eventually replaced/corrected) value. This allows the service side implementer to validate/overrule the new field value provided by a client.

The effective field value returned by the “SetHandler” is implicitly taken over by the Communication Management implementation as if the service implementer had called `Update()` explicitly with the effective value on its own. That means: An explicit `Update()` call within the “SetHandler” is superfluous as the Communication Management would update the field value with the returned value of the “SetHandler” anyways.

5.4.8.3 Ensuring existence of “SetHandler”

The existence of a registered “SetHandler” is ensured by an `ara::com` compliant implementation by returning a recoverable error: If a developer calls `OfferService()` on a skeleton implementation and had not yet registered a “SetHandler” for each of its fields, which has setter enabled, the Communication Management implementation shall return an `Error` (`ComErrc::kSetHandlerNotSet`) indicating this error in the `ara::core::Result`.

5.4.8.4 Ensuring existence of valid Field values

Since the most basic guarantee of a field is, that it has a valid value at any time, `ara::com` has to somehow ensure, that a service implementation providing a field has to provide a value **before** the service (and therefore its field) becomes visible to potential consumers, which — after subscription to the field — expect to get initial value notification event (if field is configured with notification) or a valid value on a `Get()` call (if getter is enabled for the field).

An `ara::com` Communication Management implementation needs therefore behave in the following way: If a developer calls `OfferService()` on a skeleton implementation and had not yet called `Update()` on any field, which

- has notification enabled
- or has getter enabled but not yet a “GetHandler” registered

the Communication Management implementation shall return an `Error` (`ComErrc::kFieldValueIsValid`) indicating this error in the `ara::core::Result`.

Note: The AUTOSAR meta-model supports the definition of such initial values for a field in terms of a so called `FieldSenderComSpec` of a `PPortPrototype`. So this model element should be considered by the application code calling `Update()`.

5.4.8.5 Access to current field value from Get/SetHandler

Since the underlying `field` value is only known to the middleware, the current `field` value is not accessible from the “Get/SetHandler” implementation, which are on application level. If the “Get/SetHandler” needs to read the current `field` value, the skeleton implementation must provide a `field` value replica accessible from application level.

5.4.9 Triggers

As in [Section 5.3.8](#), `Triggers` are based on `Events`, but without containing any data. Focusing on the difference from `Skeleton` side, only `Send()` is different from `Events`.

Other Event APIs aren't necessary as they relate to data which is not present for Triggers.

```
1     ara::core::Result<void> Send();
```

This will simply send out a Trigger.

No allocation is necessary due to the fact that Triggers have no data.

5.5 Data Types on Service Interface level

The following chapter describes the C++ language mapping in `ara::com` of the SI specific ("user defined") data types. "user defined" here means, that those data types aren't defined/mandated by `ara::com` API itself like e.g. `InstanceIdentifier`, `FindServiceHandle`, `ServiceHandleContainer` or any other data type defined by `ara::com` in its own namespace, but are specifically provided by the user defined SI description (IDL).

In the AUTOSAR Meta-Model ([1]) `CppImplementationDataTypes` have been introduced to support the specifics of the C++14 data type system appropriately.

5.5.1 Optional data elements

Record elements inside a `StructureImplementationDataType` can be defined as optional inside the meta-model, see [1].

This optionality is represented in the `ara::com` API by the template class `ara::core::Optional`. The serialization of such record elements is based on the Tag-Length-Value principle whereas `StructureImplementationDataTypes` without optional record elements do not have to make use of tags.

Details on how this serialization works is specified in [9].

The `ara::core::Optional` template parameter has the `ImplementationDataType` (also `ApplicationDataTypes` are possible) of the record element e.g. `uint32`.

Optional record elements can be used in structures for every SI element (e.g. Fields, Events and Methods). This optionality is defined on the SI level.

The structure in Listing 5.21 has the optional declared elements `current` and `health`. These elements are not mandatory present.

The consuming application has to check whether the optional elements contain a value or not during runtime. If an optional element contains a value or not depends on the providing application.

The providing application may set the value or not for this specific instance. The feature of optional contained elements provides forward and backward compatibility of the [SI](#) because new added record elements can just be ignored by old applications.

```
1 /**
2  * \brief Data structure with optional contained values.
3  */
4 struct BatteryState {
5     Voltage_t voltage;
6     Temperature_t temperature;
7     ara::core::Optional<Current_t> current;
8     ara::core::Optional<Health> health;
9 };
```

Listing 5.21: Definition of BatteryState

The Skeleton implementation in [Listing 5.22](#) provides the BatteryState structure defined in [Listing 5.21](#).

The implementation is aware of the optional labeled element `current` but not of the optional labeled element `health` due to a new version of the [SI](#). Therefore `health` is not set by the Skeleton implementation.

```
1 using namespace ara::com;
2
3 class BatteryStateImpl : public BatteryStateSkeleton {
4     public:
5         Future<BatteryState> GetBatteryState() {
6             // no asynchronous call for simplicity
7             ara::core::Promise<BatteryState> promise;
8
9             // fill the data structure
10            BatteryState state;
11            state.voltage = 14;
12            state.temperature = 35;
13            state.current = 0;
14            // state.health is not set and therefore it is not transmitted
15
16            promise.set_value(state);
17            auto future = promise.get_future();
18            return future;
19        }
20 }
```

Listing 5.22: Handling of optional data elements on Skeleton side

The `Proxy` in [Listing 5.23](#) consumes the `BatteryState` structure defined in [Listing 5.21](#).

The implementation is aware of both optional labeled elements `current` and `health`. Before accessing the value of the optional elements the implementation has to check whether there is really a value contained. Therefore the `optional` API provides two methods: The operator `bool` and the `has_value` method.

```
1 using namespace ara::com;
2
3 int main() {
4     // some code to acquire handle
5     // ...
6     BatteryStateProxy bms_service(handle);
7     Future<BatteryState> stateFuture = bms_service.GetBatteryState();
8     // Receive BatteryState
9     BatteryState state = stateFuture.get();
10
11    // Check the optional contained elements for presence
12    if(state.current) {
13        //Access the value of the optional element with the optional::operator*
14        if(*state.current >= MAX_CURRENT) {
15            // do something with this information
16        }
17    }
18
19    // Check with optional::has_value() method
20    if(state.health.has_value()){
21        // Access the value of the optional element with the optional::value()
22        // method
23        if(state.health.value() >= BAD_HEALTH) {
24            // do something with this information
25        }
26    }
```

Listing 5.23: Handling of optional data elements on Proxy side

5.6 Communication Groups

NOTE: Communication Groups have been set to obsolete and removed from [3], therefore it is no longer supported in future AUTOSAR release(s). This section has been left merely as reference.

5.6.1 Objective

The Communication Group (CG) is a composite service template defined by AUTOSAR. It provides a communication framework which allows to communicate peer to peer and in broadcast mode between AUTOSAR applications. The Communication Group has one server and multiple clients. The server sends messages to the clients in peer to peer and/or in broadcast mode. The clients can response to a message. The

server of a Communication Groups knows all clients connected to the Communication Group by any time. Therefore, the server can directly communicate to a certain client.

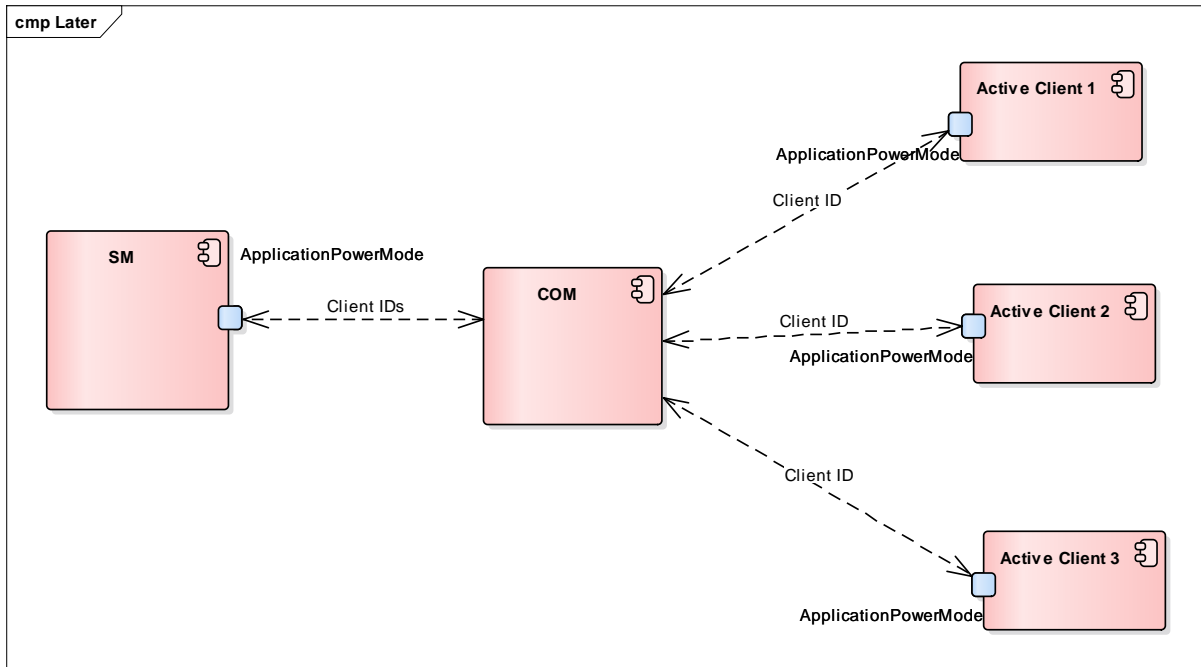


Figure 5.3: Communication Group Abstract for State Management

5.6.2 Realization

The Communication Group provides two `ServiceInterface`s.

- A dedicated Server `ServiceInterface` with the category value `COMMUNICATION_GROUP_SERVER`, implemented as `Skeleton`.
- A dedicated Client `ServiceInterface` with the category value `COMMUNICATION_GROUP_CLIENT`, implemented as `Proxy`.

A Communication Group server connects to the Communication Group using the `ServiceProxy` of the Server `ServiceInterface` and a Communication Group client connects to the Communication Group using the `ServiceSkeleton` of the Client `ServiceInterface`. See also the figure below.

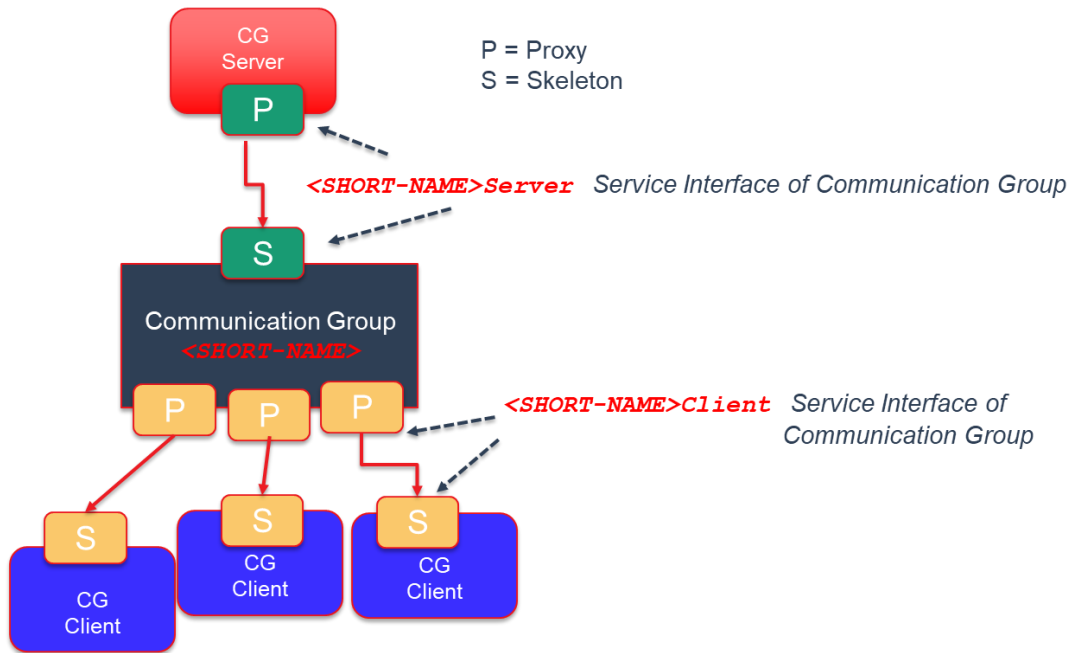


Figure 5.4: Communication Group

Note: Message types of a Communication Group are defined with the Communication Group meta-model, see also following chapters.

5.6.3 Usage Scope

The initial usage of the Communication Group has been Adaptive State Management applications. The `SWS_StateManagement` [4] specification has defined two Communication Groups, named: `PowerMode` and `DiagnosticReset`. The usage scope of the State Management Communication Groups are applications within a machine only. But from design point of view, a Communication Group supports the connection of local and remote applications (CP and AP). The figure below presents a local and remote scenarios for Communication Groups.

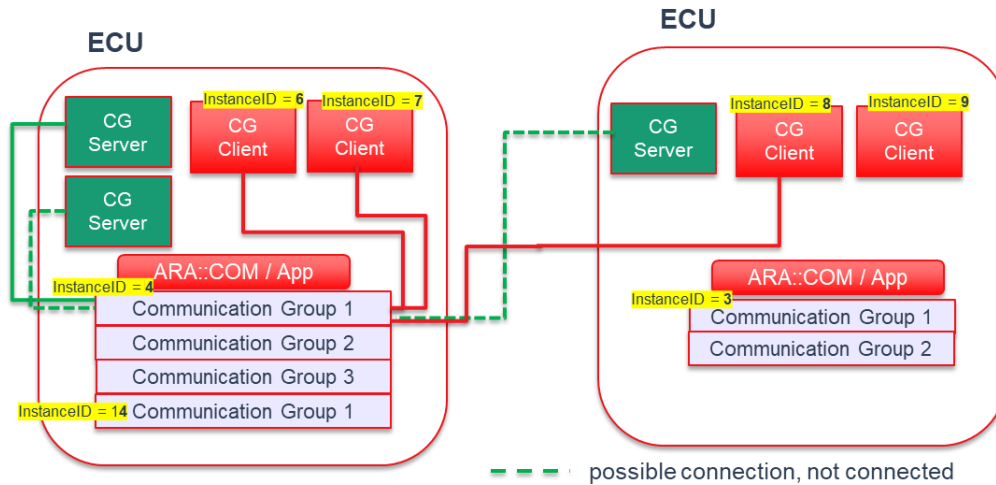


Figure 5.5: Communication Group within a Vehicle

5.6.4 Special Topics

5.6.4.1 Architecture

As before described the Communication Group contains

- only one Service Skeleton instantiation of type Communication Group server (ServiceInterface with the category value `COMMUNICATION_GROUP_SERVER`).

Note: Today there is the constraint that only one Communication server can connect to the Communication Group server ServiceInterface.

- multiple Service Proxy instantiations of type Communication Group client (ServiceInterface with the category value `COMMUNICATION_GROUP_CLIENT`).

Note: The number of Service Proxy instantiations is implementation and use case dependent. The theoretical limit (not considering for memory space requirements) is the available number of Communication client *Client_IDs*. On the other side there might be Safety requirements which force the design to support a limited known number of Communication Group clients.

5.6.4.2 Remote Connection

As mentioned before the Communication Group server or clients can be remote from the Communication Group instance. This is supported since the Communication Group is designed as composite service. Therefore also no limit in terms of AUTOSAR Classic and Adaptive platforms.

5.6.4.3 Service Versioning

Service contract versioning is supported by the Communication Group by the according `ServiceInterfaces`.

6 Tutorials

This selection of tutorials shows some minimal examples on how to use the fundamentals and features of ara::com.

6.1 Usage of Service Interfaces

The `ara::com` model elements related to both design and deployment are included in the `Manifest`. Since not all the model elements are relevant in all the development phases, the `Manifest` can be divided in different partitions:

- **Machine Manifest:** Specifies where the Adaptive AUTOSAR Software Stack is running. In the `MachineDesign` the `CommunicationSystem` structure is specified. This includes `CommunicationConnectors` and `NetworkEndpoints`.
- **SI Manifest:** Specifies the events, methods and fields that a `Service` provides.
- **Execution Manifest:** Specifies the information related to the deployment of the applications. This includes `Executables` and `Processes`. The executable references an `rootSwComponentPrototype`.
- **RPortPrototypes and PPortPrototypes:** The `rootSwComponentPrototype` has an `Application Type` that defines `PPortPrototypes` and `RPortPrototypes` and they reference the corresponding `SI`

For all the details about the manifest specification please see the [1].

Assuming that the previously mentioned partitions of the `Manifest` exist, the following sections describe the deployment of `SI` and `ServiceInstance`. Finally the most relevant aspects related with the implementation are also introduced.

6.1.1 Service Interface Deployment

The `SI Deployment` describes how the `SI` will communicate over the network. The following information must be provided:

- Reference to the network binding used in the `SI Deployment` (e.g. `SOME/IP, DDS`)
- `SI ID`
- Deployment information for the `Events, Methods and Fields`. This includes `IDs` and any `Network Binding` specific information (e.g. `Transport Protocol`)

```
<AR-PACKAGE>
  <SHORT-NAME>ServiceInterfaceDeployments</SHORT-NAME>
  <ELEMENTS>
    <SOMEIP-SERVICE-INTERFACE-DEPLOYMENT>
      <SHORT-NAME>MyInterface_SOMEIP</SHORT-NAME>
```

```

<EVENT-DEPLOYMENTS>
  <SOMEIP-EVENT-DEPLOYMENT>
    <SHORT-NAME>Counter_SOMEIP</SHORT-NAME>
    <EVENT-REF DEST="VARIABLE-DATA-PROTOTYPE"/>/ServiceInterfaces/
      myInterface/EventCounter</EVENT-REF>
    <EVENT-ID>1</EVENT-ID>
    <TRANSPORT-PROTOCOL>UDP</TRANSPORT-PROTOCOL>
  </SOMEIP-EVENT-DEPLOYMENT>
</EVENT-DEPLOYMENTS>
<SERVICE-INTERFACE-REF DEST="SERVICE-INTERFACE"/>/ServiceInterfaces/
  myInterface</SERVICE-INTERFACE-REF>
<SERVICE-INTERFACE-ID>99</SERVICE-INTERFACE-ID>
<SERVICE-INTERFACE-VERSION>
  <MAJOR-VERSION>1</MAJOR-VERSION>
  <MINOR-VERSION>0</MINOR-VERSION>
</SERVICE-INTERFACE-VERSION>
</SOMEIP-SERVICE-INTERFACE-DEPLOYMENT>
</ELEMENTS>
</AR-PACKAGE>

```

Listing 6.1: Example Service Interface Deployment

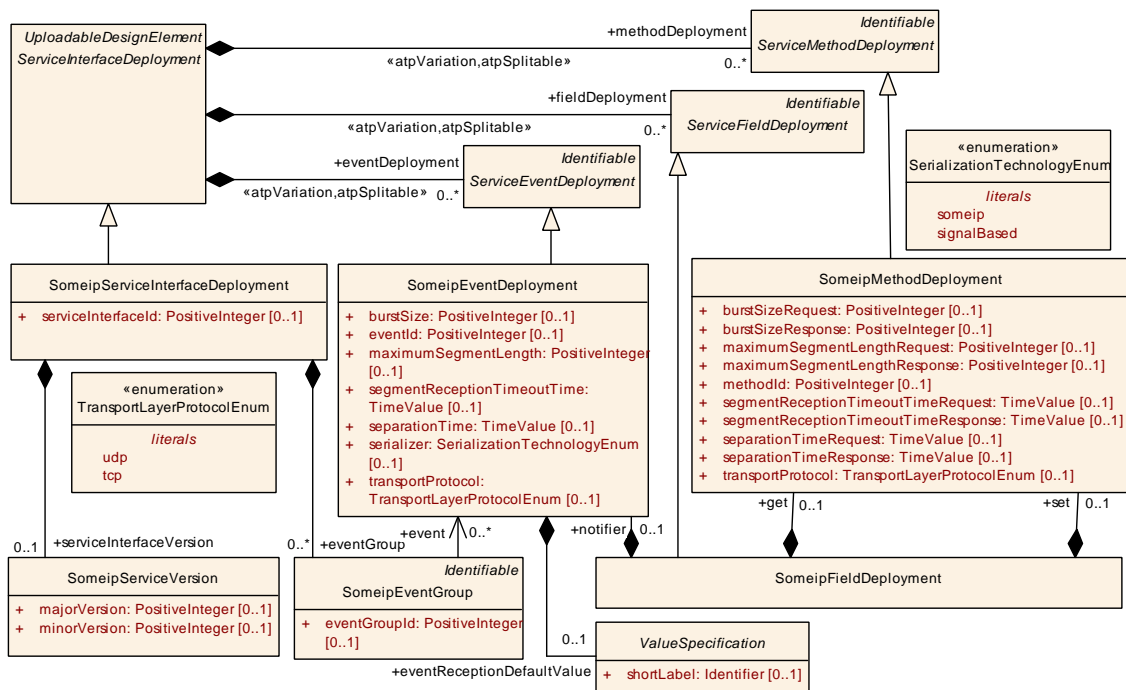


Figure 6.1: SOME/IP Service Interface Deployment

6.1.2 Service Instance Deployment

The Service Instance Deployment consists of Service Instance mapping to Application Endpoints and to Machine.

The mapping to `ApplicationEndpoints` connects a `ServiceInstance` to a `PortPrototype` and `Process`.

```

<AR-PACKAGE>
  <SHORT-NAME>ServiceInstanceToPortPrototypeMappings</SHORT-NAME>
  <ELEMENTS>
    <SERVICE-INSTANCE-TO-PORT-PROTOTYPE-MAPPING>
      <SHORT-NAME>ProvidedServiceInstance_toPort</SHORT-NAME>
      <PORT-PROTOTYPE-IREF>
        <CONTEXT-ROOT-SW-COMPONENT-PROTOTYPE-REF DEST="ROOT-SW-COMPONENT-
          PROTOTYPE"/>/Executables/myExecutable/mySwComponentPrototype</
          CONTEXT-ROOT-SW-COMPONENT-PROTOTYPE-REF>
        <TARGET-PORT-PROTOTYPE-REF DEST="P-PORT-PROTOTYPE"/>/
          SwComponentTypes/Publisher/myInterface_PPort</TARGET-PORT-
          PROTOTYPE-REF>
      </PORT-PROTOTYPE-IREF>
      <PROCESS-REF DEST="PROCESS"/>/Processes/myProcess</PROCESS-REF>
      <SERVICE-INSTANCE-REF DEST="PROVIDED-SOMEIP-SERVICE-INSTANCE"/>/
        ServiceInstances/ProvidedSomeipServiceInstance</SERVICE-INSTANCE-
        REF>
    </SERVICE-INSTANCE-TO-PORT-PROTOTYPE-MAPPING>
  </ELEMENTS>
</AR-PACKAGE>
    
```

Listing 6.2: Example Service Interface to Port Prototype Mapping

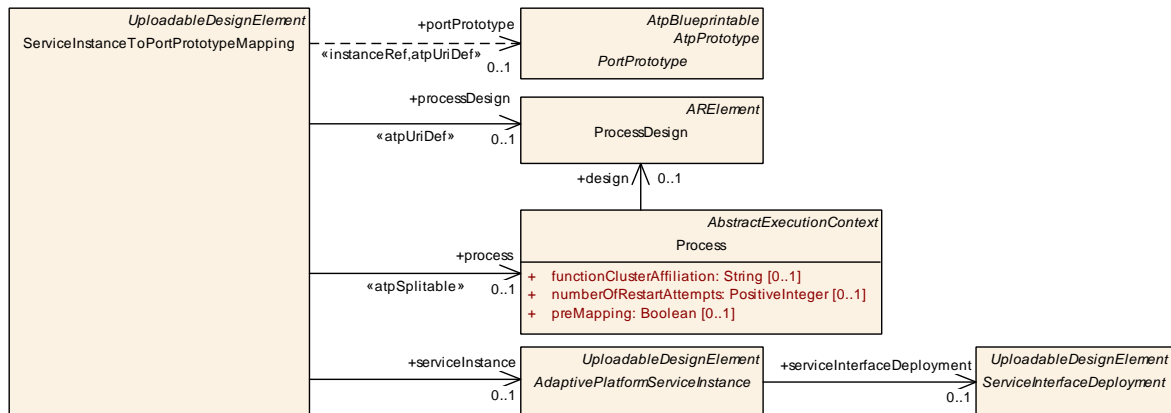


Figure 6.2: Service Instance to Port Prototype Mapping

The mapping to a `MachineDesign` connects the `ServiceInstance` to a `CommunicationConnector`.

```

<AR-PACKAGE>
  <SHORT-NAME>ServiceInstanceToMachineMappings</SHORT-NAME>
  <ELEMENTS>
    <SOMEIP-SERVICE-INSTANCE-TO-MACHINE-MAPPING>
      <SHORT-NAME>ProvidedServiceInstance_toMachine</SHORT-NAME>
      <COMMUNICATION-CONNECTOR-REF DEST="ETHERNET-COMMUNICATION-CONNECTOR">
        /MachineDesigns/publisherMachineDesign/
        publisherCommunicationConnector</COMMUNICATION-CONNECTOR-REF>
      <SERVICE-INSTANCE-REFS>
        <SERVICE-INSTANCE-REF DEST="PROVIDED-SOMEIP-SERVICE-INSTANCE"/>/
          ServiceInstances/ProvidedSomeipServiceInstance</SERVICE-INSTANCE
          -REF>
    </SOMEIP-SERVICE-INSTANCE-TO-MACHINE-MAPPING>
  </ELEMENTS>
</AR-PACKAGE>
    
```

```

</SERVICE-INSTANCE-REFS>
<UDP-PORT-REF DEST="AP-APPLICATION-ENDPOINT"/>/MachineDesigns/
  publisherMachineDesign/publisherCommunicationConnector/
  publisherEndpoint</UDP-PORT-REF>
</SOMEIP-SERVICE-INSTANCE-TO-MACHINE-MAPPING>
</ELEMENTS>
</AR-PACKAGE>

```

Listing 6.3: Example Service Instance to Machine Mapping

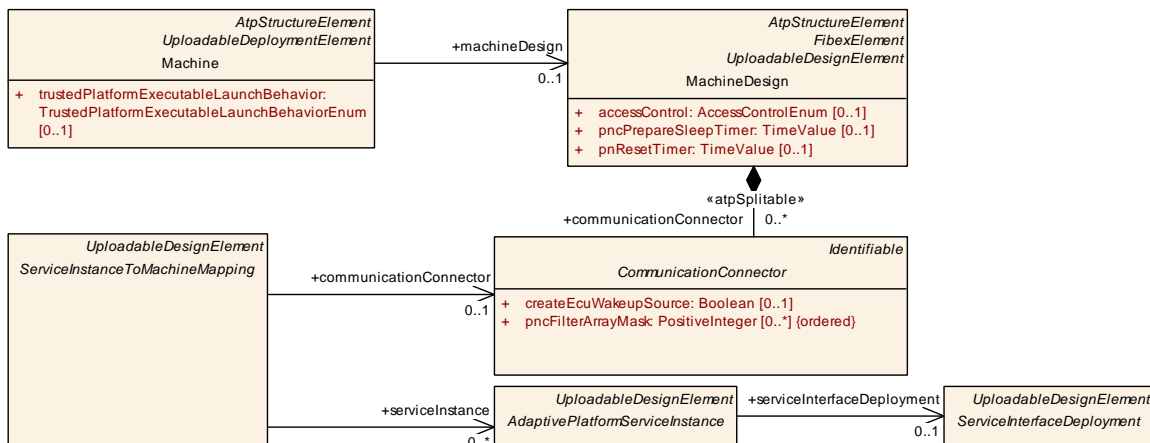


Figure 6.3: Service InstanceToMachineMapping

The `CommunicationConnector` references a `NetworkEndpoint` that includes Address, Network Mask and Gateway.

```

<AR-PACKAGE>
  <SHORT-NAME>MachineDesigns</SHORT-NAME>
  <ELEMENTS>
    <MACHINE-DESIGN>
      <SHORT-NAME>publisherMachineDesign</SHORT-NAME>
      <COMMUNICATION-CONNECTORS>
        <ETHERNET-COMMUNICATION-CONNECTOR>
          <SHORT-NAME>publisherCommunicationConnector</SHORT-NAME>
          <AP-APPLICATION-ENDPOINTS>
            <AP-APPLICATION-ENDPOINT>
              <SHORT-NAME>publisherEndpoint</SHORT-NAME>
              <TP-CONFIGURATION>
                <UDP-TP>
                  <UDP-TP-PORT>
                    <PORT-NUMBER>33222</PORT-NUMBER>
                  </UDP-TP-PORT>
                </UDP-TP>
              </TP-CONFIGURATION>
            </AP-APPLICATION-ENDPOINT>
          </AP-APPLICATION-ENDPOINTS>
          <NETWORK-ENDPOINT-REFS>
            <NETWORK-ENDPOINT-REF DEST="NETWORK-ENDPOINT"/>/system/
              theEthCluster/theEthPhysChannell/publisherNetworkEndpoint</
              NETWORK-ENDPOINT-REF>
          </NETWORK-ENDPOINT-REFS>
        </ETHERNET-COMMUNICATION-CONNECTOR>

```



```

</COMMUNICATION-CONNECTORS>
<SERVICE-DISCOVERY-CONFIGS>
  <SOMEIP-SERVICE-DISCOVERY>
    <MULTICAST-SD-IP-ADDRESS-REF DEST="NETWORK-ENDPOINT"/>/system/
      theEthCluster/theEthPhysChannel1/SDNetworkEndpoint</MULTICAST-
        SD-IP-ADDRESS-REF>
    <SOMEIP-SERVICE-DISCOVERY-PORT>30491</SOMEIP-SERVICE-DISCOVERY-
      PORT>
  </SOMEIP-SERVICE-DISCOVERY>
</SERVICE-DISCOVERY-CONFIGS>
</MACHINE-DESIGN>
</ELEMENTS>
</AR-PACKAGE>

```

Listing 6.4: Example Machine Design

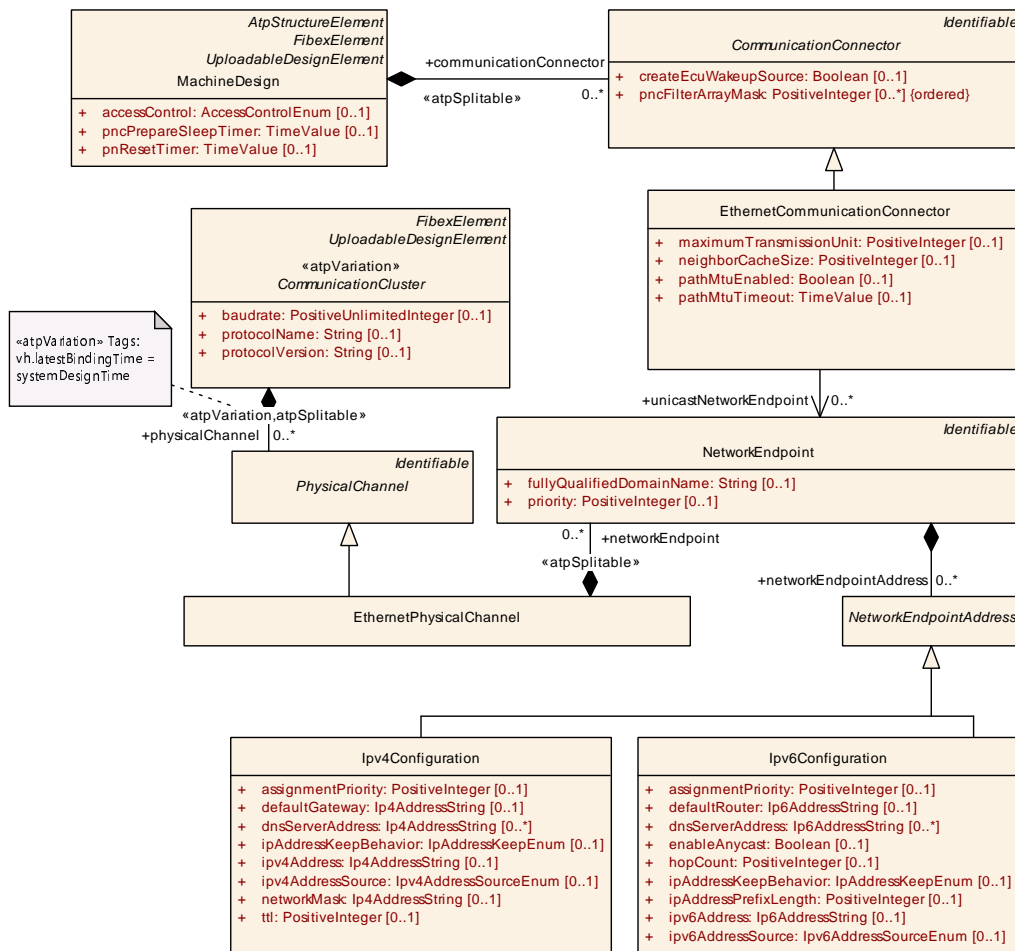


Figure 6.4: Network Connection

6.1.3 Service Implementation

Services are implemented in an application layer and are also used by the other applications. To enable the communication between both applications the Service Discovery protocol is used.




Figure 6.5: Service Discovery

The Communication Management provides a method in the Service Skeleton class to offer the service:

```
ara::core::Result<void> OfferService();
```

After the method `OfferService()` has been called, the service can be found by the applications. The Communication Management provides in the Service Proxy class a methods to find a service according to the `InstanceIdentifier` or `InstanceSpecifier`:

```
static ara::core::Result<ara::com::ServiceHandleContainer<
<ProxyClassName>::HandleType>>
FindService(ara::com::InstanceIdentifier instance);

static ara::core::Result<ara::com::ServiceHandleContainer<
<ProxyClassName>::HandleType>>
FindService(ara::core::InstanceSpecifier instance);
```

These methods return a `ServiceHandleContainer` that can have different number of elements:

- Empty: No Service Instance was found
- One: A single Service instance was found
- Several: Multiple Service instances were found

The example in [5.5](#) (lines 52 - 60) shows how to create a `Service Proxy` using `FindService()`.

A `Service Handler` gives access to all the events, methods and fields of the service. For the events, the application has to subscribe to the events using the corresponding method setting also the size of the cache for this event :

```
ara::core::Result<void> Event::Subscribe(size_t maxSampleCount);
```

To unsubscribe to an event the `Unsubscribe` method must be used:

```
void Event::Unsubscribe();
```

The example in [5.5](#) (lines 62 - 64) shows how to subscribe for an event.

The application code for the `Service Proxy` and `Skeleton` is generated according to the `SI` defined in the `Manifest`. An example of the generated code for the `Proxy Class` can be found in the [5.2](#) and for the `Skeleton Class` in [5.12](#)

6.2 Usage of InstanceSpecifier

`InstanceSpecifier` is a core concept defined in [\[2\]](#), representing a "/"-separated list of `shortNames` of model elements conforming an absolute path to an specific model element. In less formal terms, Instance Specifiers bridge Adaptive Platform models and applications, allowing application code to unequivocally reference resource instances defined by the system model.

The instances referenced by Instance Specifiers may be of many different kinds: provided or required service instances, key/value or file stores, or cryptographic resources, to name a few.

In the C++ language binding of the Adaptive Platform, the `InstanceSpecifier` class has several common traits:

- It is **not** default constructible.
- It is copiable and movable.
- It is comparable against `StringView` and other `InstanceSpecifier` objects.
- It is explicitly convertible to `StringView`.

See listing [Listing 4.2](#) in chapter [Section 4.8.1](#) for more detailed information.

Considering these characteristics, the only way to create a new `InstanceSpecifier` object that's not a copy or a move of an existing one is to do so from a `StringView` object.

In practice, the contents of such `StringView` are syntactically but not semantically checked upon construction. This means that construction succeeds as long as the source `StringView` object contains a "/"-separated list of names conformed only of valid characters. Whether the path described by such list is actually valid in the model

from which the application derives, is something that will be checked by the different functional clusters when attempting to access or instantiate resources pointed by the `InstanceSpecifier` in question.

The following examples show the way in which, according to [3], `InstanceSpecifiers` might be used to instantiate and access services via skeleton and proxy classes.

Instance specifiers point to instances of port prototypes associated with a service. Therefore multiple instance specifiers for each `PortPrototype` can be created. In the examples below, `SwComponentInstance_0` and `SwComponentInstance_1` are 2 instantiations of the the same `SwComponentPrototype`, containing the `RPortPrototype` `RPort_3`:

```

1 #include "ara/core/instance_specifier.h"
2 #include "ara/com/sample/tire_skeleton.h"
3
4 class TireSkeletonImplementation : public ara::com::sample::skeleton::
    TireSkeleton
5 {
6     using TireSkeleton::TireSkeleton;
7
8     // Implement service interface methods, if any
9 };
10
11 int main()
12 {
13     const ara::core::InstanceSpecifier tire0_Instance{"/ServerExe/
    RootSWCP_0/Comp_Lvl1/Comp_Lvl2/SwComponentInstance_0/RPort_3"};
14     const ara::core::InstanceSpecifier tire1_Instance{"/ServerExe/
    RootSWCP_0/Comp_Lvl1/Comp_Lvl2/SwComponentInstance_1/RPort_3"};
15
16     TireSkeletonImplementation tire0(tire0_Instance);
17     TireSkeletonImplementation tire1(tire1_Instance);
18
19     // Sleep while Skeleton instances run, process requests, etc.
20
21     return 0;
22 }
    
```

Listing 6.5: Example Usage of Instance Specifiers with Skeletons

```

1 #include "ara/core/instance_specifier.h"
2 #include "ara/com/sample/tire_proxy.h"
3
4 int main()
5 {
6     using Proxy = ara::com::sample::proxy::TireProxy;
7
8     const ara::core::InstanceSpecifier tire0_Instance{"/ClientExe/
    RootSWCP_0/Comp_Lvl1/Comp_Lvl2/SwComponentInstance_0/RPort_3"};
9     const ara::core::InstanceSpecifier tire1_Instance{"/ClientExe/
    RootSWCP_0/Comp_Lvl1/Comp_Lvl2/SwComponentInstance_1/RPort_3"};
10
11     auto tire0_handles = Proxy::FindService(tire0_Instance).ValueOrThrow();
12     auto tire1_handles = Proxy::FindService(tire1_Instance).ValueOrThrow();
13
    
```

```

14     Proxy tire0(tire0_handles[0]);
15     Proxy tire1(tire1_handles[0]);
16
17     // Call methods, subscribe to events, etc.
18
19     return 0;
20 }
    
```

Listing 6.6: Example Usage of Instance Specifiers with Proxies

In these examples the AUTOSAR Adaptive implementation manages process-specific manifests with isolated contexts. These can be dictated on process startup via e.g. command line arguments, environment variables, working directory contents or any other implementation-specific means.

6.2.1 Modeling and configuration/mapping over Manifest from user perspective

The `InstanceSpecifier` used for finding a service maps to the particular instance of the port associated with that service.

```

1 #include "ara/core/instance_specifier.h"
2 #include "ara/com/sample/radar_proxy.h"
3
4
5 int main() {
6     using Proxy = ara::com::sample::proxy::radarProxy;
7     // ...
8     ara::core::InstanceSpecifier portSpecifier{"/fusionExe/fusion/
9     radar_RPort"};
10    auto res = Proxy::FindService(portSpecifier);
11    //.....
12 }
    
```

Listing 6.7: CPP Example Usage with FindService

In the Application Design, the Executable node specifies its `RootSwComponentPrototype`. In turn the `SwComponentPrototype` defines one or more `PortPrototype`.

```

<EXECUTABLE>
  <SHORT-NAME>fusionExe</SHORT-NAME>
  <CATEGORY>APPLICATION_LEVEL</CATEGORY>
  <ROOT-SW-COMPONENT-PROTOTYPE>
    <SHORT-NAME>fusion</SHORT-NAME>
    <APPLICATION-TYPE-TREF DEST="ADAPTIVE-APPLICATION-SW-COMPONENT-TYPE"/>
    apd/da/fusion</APPLICATION-TYPE-TREF>
  </ROOT-SW-COMPONENT-PROTOTYPE>
</EXECUTABLE>

<ADAPTIVE-APPLICATION-SW-COMPONENT-TYPE>
  <SHORT-NAME>fusion</SHORT-NAME>
  <PORTS>
    <R-PORT-PROTOTYPE>
      <SHORT-NAME>radar_RPort</SHORT-NAME>
    
```

```

        <REQUIRED-INTERFACE-TREF DEST="SERVICE-INTERFACE"/>ServiceInterfaces/
            radar</REQUIRED-INTERFACE-TREF>
        </R-PORT-PROTOTYPE>
    </PORTS>
</ADAPTIVE-APPLICATION-SW-COMPONENT-TYPE>

<SOMEIP-SERVICE-INTERFACE-DEPLOYMENT>
    <SHORT-NAME>radar_Someip</SHORT-NAME>
    <!-- ..... -->
</SOMEIP-SERVICE-INTERFACE-DEPLOYMENT>
    
```

Listing 6.8: Path towards port instance

See [Figure 6.1](#) for the SOME/IP Service Interface Deployment.

The mapping between a `RequiredServiceInstance` and an `InstanceSpecifier` is done via the *Service Instance Manifest*. In the *Service Instance Manifest* the `ServiceInstanceToPortPrototypeMapping` defines which *Service Instance* is associated with a certain port inside a specific `RootSwComponentPrototype`. The `RequiredServiceInstance` specifies the `InstanceId` as `RequiredServiceInstanceId`, in the example below this value is 19.

```

<?xml version="1.0" encoding="UTF-8"?>
<AUTOSAR xmlns="http://autosar.org/schema/r4.0"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="http://autosar.org/schema/r4.0_AUTOSAR_00053.xsd">
    <AR-PACKAGES>
        <AR-PACKAGE>
            <SHORT-NAME>apd</SHORT-NAME>
            <AR-PACKAGES>
                <AR-PACKAGE>
                    <SHORT-NAME>da</SHORT-NAME>
                    <AR-PACKAGES>
                        <AR-PACKAGE>
                            <SHORT-NAME>instance</SHORT-NAME>
                            <ELEMENTS>
                                <SOMEIP-SERVICE-INSTANCE-TO-MACHINE-MAPPING>
                                    <SHORT-NAME>radar_RequiredServiceInstance_toMachine</SHORT-NAME>
                                    <COMMUNICATION-CONNECTOR-REF DEST="ETHERNET-COMMUNICATION-CONNECTOR"/>MachineDesigns/fusionMachineDesign/fusionCommunicationConnector</COMMUNICATION-CONNECTOR-REF>
                                    <SERVICE-INSTANCE-REFS>
                                        <SERVICE-INSTANCE-REF DEST="REQUIRED-SOMEIP-SERVICE-INSTANCE"/>apd/da/instance/radar_RequiredSomeipServiceInstance</SERVICE-INSTANCE-REF>
                                    </SERVICE-INSTANCE-REFS>
                                    <UDP-PORT-REF DEST="AP-APPLICATION-ENDPOINT"/>MachineDesigns/fusionMachineDesign/fusionCommunicationConnector/fusionEndpoint</UDP-PORT-REF>
                                </SOMEIP-SERVICE-INSTANCE-TO-MACHINE-MAPPING>
                            </ELEMENTS>
                        </AR-PACKAGE>
                    </AR-PACKAGES>
                </AR-PACKAGE>
            </AR-PACKAGES>
        </AR-PACKAGE>
    </AR-PACKAGES>
</AUTOSAR>
    
```

```

</SOMEIP-SERVICE-INSTANCE-TO-MACHINE-MAPPING>
<SERVICE-INSTANCE-TO-PORT-PROTOTYPE-MAPPING>
  <SHORT-NAME>radar_RequiredServiceInstance_toPort</SHORT-
    NAME>
  <PORT-PROTOTYPE-IREF>
    <CONTEXT-ROOT-SW-COMPONENT-PROTOTYPE-REF DEST="ROOT-SW-
      COMPONENT-PROTOTYPE"/>Executables/fusionExe/fusion</
      CONTEXT-ROOT-SW-COMPONENT-PROTOTYPE-REF>
    <TARGET-PORT-PROTOTYPE-REF DEST="R-PORT-PROTOTYPE"/>/
      SwComponentTypes/fusion/radar_RPort</TARGET-PORT-
      PROTOTYPE-REF>
  </PORT-PROTOTYPE-IREF>
  <PROCESS-REF DEST="PROCESS"/>/Processes/fusion_instancel</
    PROCESS-REF>
  <SERVICE-INSTANCE-REF DEST="REQUIRED-SOMEIP-SERVICE-
    INSTANCE"/>/apd/da/instance/
    radar_RequiredSomeipServiceInstance</SERVICE-INSTANCE-
    REF>
</SERVICE-INSTANCE-TO-PORT-PROTOTYPE-MAPPING>

<REQUIRED-SOMEIP-SERVICE-INSTANCE>
  <SHORT-NAME>radar_RequiredSomeipServiceInstance</SHORT-
    NAME>
  <SERVICE-INTERFACE-DEPLOYMENT-REF DEST="SOMEIP-SERVICE-
    INTERFACE-DEPLOYMENT"/>/ServiceInterfaceDeployments/
    radar_Someip</SERVICE-INTERFACE-DEPLOYMENT-REF>
  <!-- ..... -->
  <REQUIRED-MINOR-VERSION>0</REQUIRED-MINOR-VERSION>
  <REQUIRED-SERVICE-INSTANCE-ID>19</REQUIRED-SERVICE-
    INSTANCE-ID>
</REQUIRED-SOMEIP-SERVICE-INSTANCE>
</ELEMENTS>
</AR-PACKAGE>
</AR-PACKAGES>
</AR-PACKAGE>
</AR-PACKAGES>
</AR-PACKAGE>
<AR-PACKAGE>
  <SHORT-NAME>MachineDesigns</SHORT-NAME>
  <ELEMENTS>
    <MACHINE-DESIGN>
      <SHORT-NAME>fusionMachineDesign</SHORT-NAME>
      <COMMUNICATION-CONNECTORS>
        <ETHERNET-COMMUNICATION-CONNECTOR>
          <SHORT-NAME>fusionCommunicationConnector</SHORT-NAME>
          <AP-APPLICATION-ENDPOINTS>
            <AP-APPLICATION-ENDPOINT>
              <SHORT-NAME>fusionEndpoint</SHORT-NAME>
              <TP-CONFIGURATION>
                <UDP-TP>
                  <UDP-TP-PORT>
                    <PORT-NUMBER>33111</PORT-NUMBER>
                  </UDP-TP-PORT>
                </UDP-TP>
              </TP-CONFIGURATION>
            </AP-APPLICATION-ENDPOINT>
          </AP-APPLICATION-ENDPOINTS>
        </ETHERNET-COMMUNICATION-CONNECTOR>
      </COMMUNICATION-CONNECTORS>
    </MACHINE-DESIGN>
  </ELEMENTS>
</AR-PACKAGE>

```

```

</AP-APPLICATION-ENDPOINTS>
<NETWORK-ENDPOINT-REFS>
  <NETWORK-ENDPOINT-REF DEST="NETWORK-ENDPOINT">/system/
    theEthCluster/theEthPhysChannel1/
      publisherNetworkEndpoint</NETWORK-ENDPOINT-REF>
</NETWORK-ENDPOINT-REFS>
</ETHERNET-COMMUNICATION-CONNECTOR>
</COMMUNICATION-CONNECTORS>
<SERVICE-DISCOVERY-CONFIGS>
  <SOMEIP-SERVICE-DISCOVERY>
    <MULTICAST-SD-IP-ADDRESS-REF DEST="NETWORK-ENDPOINT">/system/
      theEthCluster/theEthPhysChannel1/SDNetworkEndpoint</
        MULTICAST-SD-IP-ADDRESS-REF>
    <SOMEIP-SERVICE-DISCOVERY-PORT>30491</SOMEIP-SERVICE-
      DISCOVERY-PORT>
  </SOMEIP-SERVICE-DISCOVERY>
</SERVICE-DISCOVERY-CONFIGS>
</MACHINE-DESIGN>
</ELEMENTS>
</AR-PACKAGE>

<AR-PACKAGE>
  <SHORT-NAME>system</SHORT-NAME>
  <ELEMENTS>
    <ETHERNET-CLUSTER>
      <SHORT-NAME>theEthCluster</SHORT-NAME>
      <ETHERNET-CLUSTER-VARIANTS>
        <ETHERNET-CLUSTER-CONDITIONAL>
          <PHYSICAL-CHANNELS>
            <ETHERNET-PHYSICAL-CHANNEL>
              <SHORT-NAME>theEthPhysChannel1</SHORT-NAME>
              <NETWORK-ENDPOINTS>
                <NETWORK-ENDPOINT>
                  <SHORT-NAME>publisherNetworkEndpoint</SHORT-NAME>
                </NETWORK-ENDPOINT>
                <NETWORK-ENDPOINT>
                  <SHORT-NAME>SDNetworkEndpoint</SHORT-NAME>
                </NETWORK-ENDPOINT>
              </NETWORK-ENDPOINTS>
            </ETHERNET-PHYSICAL-CHANNEL>
          </PHYSICAL-CHANNELS>
        </ETHERNET-CLUSTER-CONDITIONAL>
      </ETHERNET-CLUSTER-VARIANTS>
    </ETHERNET-CLUSTER>
  </ELEMENTS>
</AR-PACKAGE>

<AR-PACKAGE>
  <SHORT-NAME>Executables</SHORT-NAME>
  <ELEMENTS>
    <EXECUTABLE>
      <SHORT-NAME>fusionExe</SHORT-NAME>
      <ROOT-SW-COMPONENT-PROTOTYPE>
        <SHORT-NAME>fusion</SHORT-NAME>
      </ROOT-SW-COMPONENT-PROTOTYPE>
    </EXECUTABLE>
  </ELEMENTS>
</AR-PACKAGE>

```



```

    </ELEMENTS>
  </AR-PACKAGE>

  <AR-PACKAGE>
    <SHORT-NAME>SwComponentTypes</SHORT-NAME>
    <ELEMENTS>
      <ADAPTIVE-APPLICATION-SW-COMPONENT-TYPE>
        <SHORT-NAME>fusion</SHORT-NAME>
        <PORTS>
          <R-PORT-PROTOTYPE>
            <SHORT-NAME>radar_RPort</SHORT-NAME>
          </R-PORT-PROTOTYPE>
        </PORTS>
      </ADAPTIVE-APPLICATION-SW-COMPONENT-TYPE>
    </ELEMENTS>
  </AR-PACKAGE>

  <AR-PACKAGE>
    <SHORT-NAME>Processes</SHORT-NAME>
    <ELEMENTS>
      <PROCESS>
        <SHORT-NAME>fusion_instancel</SHORT-NAME>
      </PROCESS>
    </ELEMENTS>
  </AR-PACKAGE>
  <AR-PACKAGE>

    <SHORT-NAME>ServiceInterfaceDeployments</SHORT-NAME>
    <ELEMENTS>
      <SOMEIP-SERVICE-INTERFACE-DEPLOYMENT>
        <SHORT-NAME>radar_Someip</SHORT-NAME>
      </SOMEIP-SERVICE-INTERFACE-DEPLOYMENT>
    </ELEMENTS>
  </AR-PACKAGE>

</AR-PACKAGES>
</AUTOSAR>

```

Listing 6.9: Service Instance Manifest specification of RequiredServiceInstance

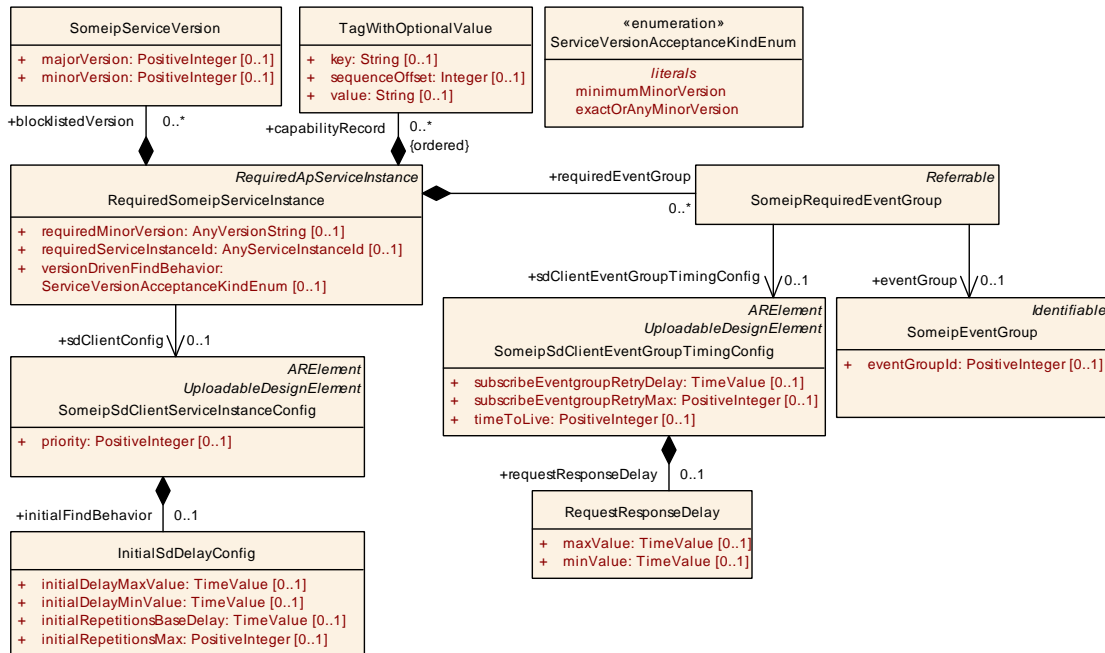


Figure 6.6: SOME/IP Required Service Instance Deployment

Note: An executable can be started multiple times. Each of the processes is started with a different Service Instance Manifest, therefore the mapping between an `InstanceSpecifier` and a `RequiredServiceInstance` will be different between the separate processes.

6.2.2 Instance IDs only for provided Services

The `InstanceSpecifier` and the `InstanceIdentifier` can be used to uniquely identify a provided service. This means that clients know which particular service instance they are communicating with. This information is lacking for provided services. Clients can't be uniquely identified with an `InstanceIdentifier`, therefore the server can't know for sure with which client it communicates with. For most cases this is not a problem, however we envision that for safety this can be a problem. For these cases we recommend using the `E2E` parameter `dataID` of the method `E2E_check` (see [PRS_E2E_00323] of [10]).

6.3 Usage in context of MultiBinding

The following section revamps the example shown in 6.2.1, depicting how multiple network bindings for a single `PortPrototype` can be defined in the Application and Instance Manifests without altering the Adaptive Application's source code. In this case, we will add an additional DDS-based instantiation of the `/apd/da/radar SI`.

To begin with, the Application Manifest is extended with an `DdsServiceInterfaceDeployment`, named `radar_Dds`, portraying DDS-specific deployment elements of the SI:

```
<?xml version="1.0" encoding="UTF-8"?>
<AUTOSAR xmlns="http://autosar.org/schema/r4.0" xmlns:xsi="http://www.w3.
  org/2001/XMLSchema-instance" xsi:schemaLocation="http://autosar.org/
  schema/r4.0_AUTOSAR_00053.xsd">
  <AR-PACKAGES>
    <AR-PACKAGE>
      <SHORT-NAME>apd</SHORT-NAME>
      <AR-PACKAGES>
        <AR-PACKAGE>
          <SHORT-NAME>da</SHORT-NAME>
          <ELEMENTS>
            <EXECUTABLE>
              <SHORT-NAME>fusionExe</SHORT-NAME>
              <CATEGORY>APPLICATION_LEVEL</CATEGORY>
              <ROOT-SW-COMPONENT-PROTOTYPE>
                <SHORT-NAME>fusion</SHORT-NAME>
                <APPLICATION-TYPE-TREF DEST="ADAPTIVE-APPLICATION-SW-
                  COMPONENT-TYPE"/>/apd/da/fusion</APPLICATION-TYPE-TREF>
              </ROOT-SW-COMPONENT-PROTOTYPE>
            </EXECUTABLE>

            <ADAPTIVE-APPLICATION-SW-COMPONENT-TYPE>
              <SHORT-NAME>fusion</SHORT-NAME>
              <PORTS>
                <P-PORT-PROTOTYPE>
                  <SHORT-NAME>radar_PPort</SHORT-NAME>
                  <PROVIDED-INTERFACE-TREF DEST="SERVICE-INTERFACE"/>/
                    ServiceInterfaces/radar</PROVIDED-INTERFACE-TREF>
                </P-PORT-PROTOTYPE>
              </PORTS>
            </ADAPTIVE-APPLICATION-SW-COMPONENT-TYPE>

            <SOMEIP-SERVICE-INTERFACE-DEPLOYMENT>
              <SHORT-NAME>radar_Someip</SHORT-NAME>
              <!-- ..... -->
            </SOMEIP-SERVICE-INTERFACE-DEPLOYMENT>

            <DDS-SERVICE-INTERFACE-DEPLOYMENT>
              <SHORT-NAME>radar_Dds</SHORT-NAME>
              <!-- ..... -->
            </DDS-SERVICE-INTERFACE-DEPLOYMENT>
          </ELEMENTS>
        </AR-PACKAGE>
      </AR-PACKAGES>
    </AR-PACKAGE>
    <AR-PACKAGE>
      <SHORT-NAME>ServiceInterfaces</SHORT-NAME>
      <ELEMENTS>
        <SERVICE-INTERFACE>
          <SHORT-NAME>radar</SHORT-NAME>
        </SERVICE-INTERFACE>
      </ELEMENTS>
    </AR-PACKAGE>
  </AR-PACKAGES>
</AUTOSAR>
```

```

</AR-PACKAGE>
</AR-PACKAGES>
</AUTOSAR>

```

Listing 6.10: Path towards port instance

See Figure 6.1 for the SOME/IP Service Interface Deployment.

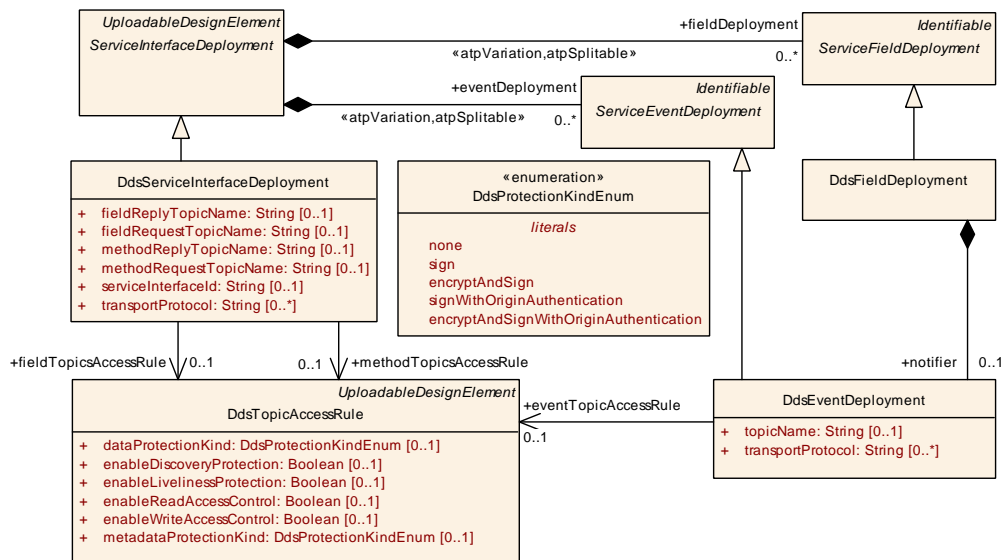


Figure 6.7: DDS Service Interface Deployment

In the Instance Manifest, separate sets of ServiceInstanceToMachineMapping, ServiceInstanceToPortPrototypeMapping and ProvidedServiceInstance are listed, each defining SOME/IP and DDS-specific deployment parameters.

```

<?xml version="1.0" encoding="UTF-8"?>
<AUTOSAR xmlns="http://autosar.org/schema/r4.0" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:schemaLocation="http://autosar.org/schema/r4.0_AUTOSAR_00053.xsd">
  <AR-PACKAGES>
    <AR-PACKAGE>
      <SHORT-NAME>apd</SHORT-NAME>
      <AR-PACKAGES>
        <AR-PACKAGE>
          <SHORT-NAME>da</SHORT-NAME>
          <AR-PACKAGES>
            <AR-PACKAGE>
              <SHORT-NAME>instance</SHORT-NAME>
              <ELEMENTS>
                <!-- For SOME/IP -->
                <SOMEIP-SERVICE-INSTANCE-TO-MACHINE-MAPPING>
                  <SHORT-NAME>
                    radar_ProvidedServiceInstance_toMachine_Someip</SHORT-NAME>

```

```

<COMMUNICATION-CONNECTOR-REF DEST="ETHERNET-
COMMUNICATION-CONNECTOR"/>/MachineDesigns/
fusionMachineDesign/fusionCommunicationConnector</
COMMUNICATION-CONNECTOR-REF>
<SERVICE-INSTANCE-REFS>
  <SERVICE-INSTANCE-REF DEST="PROVIDED-SOMEIP-SERVICE-
  INSTANCE"/>/apd/da/instance/
  radar_ProvidedSomeipServiceInstance</SERVICE-
  INSTANCE-REF>
</SERVICE-INSTANCE-REFS>
<UDP-PORT-REF DEST="AP-APPLICATION-ENDPOINT"/>/
MachineDesigns/fusionMachineDesign/
fusionCommunicationConnector/fusionEndpoint</UDP-
PORT-REF>
</SOMEIP-SERVICE-INSTANCE-TO-MACHINE-MAPPING>

<SERVICE-INSTANCE-TO-PORT-PROTOTYPE-MAPPING>
  <SHORT-NAME>radar_ProvidedServiceInstance_toPort_Someip
  </SHORT-NAME>
  <PORT-PROTOTYPE-IREF>
    <CONTEXT-ROOT-SW-COMPONENT-PROTOTYPE-REF DEST="ROOT-
    SW-COMPONENT-PROTOTYPE"/>/Executables/fusionExe/
    fusion</CONTEXT-ROOT-SW-COMPONENT-PROTOTYPE-REF>
    <TARGET-PORT-PROTOTYPE-REF DEST="P-PORT-PROTOTYPE"/>/
    SwComponentTypes/fusion/radar_PPort</TARGET-PORT-
    PROTOTYPE-REF>
  </PORT-PROTOTYPE-IREF>
  <PROCESS-REF DEST="PROCESS"/>/Processes/fusion_instancel
  </PROCESS-REF>
  <SERVICE-INSTANCE-REF DEST="PROVIDED-SOMEIP-SERVICE-
  INSTANCE"/>/apd/da/instance/
  radar_ProvidedSomeipServiceInstance</SERVICE-
  INSTANCE-REF>
</SERVICE-INSTANCE-TO-PORT-PROTOTYPE-MAPPING>

<PROVIDED-SOMEIP-SERVICE-INSTANCE>
  <SHORT-NAME>radar_ProvidedSomeipServiceInstance</SHORT-
  NAME>
  <SERVICE-INTERFACE-DEPLOYMENT-REF DEST="SOMEIP-SERVICE-
  INTERFACE-DEPLOYMENT"/>/ServiceInterfaceDeployments/
  radar_Someip</SERVICE-INTERFACE-DEPLOYMENT-REF>
  <!-- ..... -->
  <SERVICE-INSTANCE-ID>19</SERVICE-INSTANCE-ID>
</PROVIDED-SOMEIP-SERVICE-INSTANCE>

<!-- For DDS -->
<DDS-SERVICE-INSTANCE-TO-MACHINE-MAPPING>
  <SHORT-NAME>radar_ProvidedServiceInstance_toMachine_Dds
  </SHORT-NAME>
  <COMMUNICATION-CONNECTOR-REF DEST="ETHERNET-
  COMMUNICATION-CONNECTOR"/>/MachineDesigns/
  fusionMachineDesign/fusionCommunicationConnector</
  COMMUNICATION-CONNECTOR-REF>
  <SERVICE-INSTANCE-REFS>

```

```

        <SERVICE-INSTANCE-REF DEST="DDS-PROVIDED-SERVICE-
            INSTANCE"/>/apd/da/instance/
            radar_ProvidedDdsServiceInstance</SERVICE-INSTANCE
            -REF>
    </SERVICE-INSTANCE-REFS>
</DDS-SERVICE-INSTANCE-TO-MACHINE-MAPPING>

<SERVICE-INSTANCE-TO-PORT-PROTOTYPE-MAPPING>
    <SHORT-NAME>radar_ProvidedServiceInstance_toPort_Dds</
        SHORT-NAME>
    <PORT-PROTOTYPE-IREF>
        <CONTEXT-ROOT-SW-COMPONENT-PROTOTYPE-REF DEST="ROOT-
            SW-COMPONENT-PROTOTYPE"/>Executables/fusionExe/
            fusion</CONTEXT-ROOT-SW-COMPONENT-PROTOTYPE-REF>
        <TARGET-PORT-PROTOTYPE-REF DEST="P-PORT-PROTOTYPE"/>/
            SwComponentTypes/fusion/radar_PPort</TARGET-PORT-
            PROTOTYPE-REF>
    </PORT-PROTOTYPE-IREF>
    <PROCESS-REF DEST="PROCESS"/>/Processes/fusion_instancel
    </PROCESS-REF>
    <SERVICE-INSTANCE-REF DEST="DDS-PROVIDED-SERVICE-
        INSTANCE"/>/apd/da/instance/
        radar_ProvidedDdsServiceInstance</SERVICE-INSTANCE-
        REF>
</SERVICE-INSTANCE-TO-PORT-PROTOTYPE-MAPPING>

<DDS-PROVIDED-SERVICE-INSTANCE>
    <SHORT-NAME>radar_ProvidedDdsServiceInstance</SHORT-
        NAME>
    <SERVICE-INTERFACE-DEPLOYMENT-REF DEST="DDS-SERVICE-
        INTERFACE-DEPLOYMENT"/>/ServiceInterfaceDeployments/
        radar_Dds</SERVICE-INTERFACE-DEPLOYMENT-REF>
    <!-- ..... -->
    <SERVICE-INSTANCE-ID>19</SERVICE-INSTANCE-ID>
</DDS-PROVIDED-SERVICE-INSTANCE>
</ELEMENTS>
</AR-PACKAGE>
</AR-PACKAGES>
</AR-PACKAGE>
</AR-PACKAGES>
</AR-PACKAGE>
<AR-PACKAGE>
    <SHORT-NAME>MachineDesigns</SHORT-NAME>
    <ELEMENTS>
        <MACHINE-DESIGN>
            <SHORT-NAME>fusionMachineDesign</SHORT-NAME>
            <COMMUNICATION-CONNECTORS>
                <ETHERNET-COMMUNICATION-CONNECTOR>
                    <SHORT-NAME>fusionCommunicationConnector</SHORT-NAME>
                    <AP-APPLICATION-ENDPOINTS>
                        <AP-APPLICATION-ENDPOINT>
                            <SHORT-NAME>fusionEndpoint</SHORT-NAME>
                            <TP-CONFIGURATION>
                                <UDP-TP>
                                    <UDP-TP-PORT>
                                        <PORT-NUMBER>33111</PORT-NUMBER>

```

```

        </UDP-TP-PORT>
    </UDP-TP>
</TP-CONFIGURATION>
</AP-APPLICATION-ENDPOINT>
</AP-APPLICATION-ENDPOINTS>
<NETWORK-ENDPOINT-REFS>
    <NETWORK-ENDPOINT-REF DEST="NETWORK-ENDPOINT">/system/
        theEthCluster/theEthPhysChannel1/
        publisherNetworkEndpoint</NETWORK-ENDPOINT-REF>
</NETWORK-ENDPOINT-REFS>
</ETHERNET-COMMUNICATION-CONNECTOR>
</COMMUNICATION-CONNECTORS>
<SERVICE-DISCOVERY-CONFIGS>
    <SOMEIP-SERVICE-DISCOVERY>
        <MULTICAST-SD-IP-ADDRESS-REF DEST="NETWORK-ENDPOINT">/
            system/theEthCluster/theEthPhysChannel1/
            SDNetworkEndpoint</MULTICAST-SD-IP-ADDRESS-REF>
        <SOMEIP-SERVICE-DISCOVERY-PORT>30491</SOMEIP-SERVICE-
            DISCOVERY-PORT>
    </SOMEIP-SERVICE-DISCOVERY>
</SERVICE-DISCOVERY-CONFIGS>
</MACHINE-DESIGN>
</ELEMENTS>
</AR-PACKAGE>
<AR-PACKAGE>
    <SHORT-NAME>Executables</SHORT-NAME>
    <ELEMENTS>
        <EXECUTABLE>
            <SHORT-NAME>fusionExe</SHORT-NAME>
            <ROOT-SW-COMPONENT-PROTOTYPE>
                <SHORT-NAME>fusion</SHORT-NAME>
            </ROOT-SW-COMPONENT-PROTOTYPE>
        </EXECUTABLE>
    </ELEMENTS>
</AR-PACKAGE>

<AR-PACKAGE>
    <SHORT-NAME>SwComponentTypes</SHORT-NAME>
    <ELEMENTS>
        <ADAPTIVE-APPLICATION-SW-COMPONENT-TYPE>
            <SHORT-NAME>fusion</SHORT-NAME>
            <PORTS>
                <P-PORT-PROTOTYPE>
                    <SHORT-NAME>radar_PPort</SHORT-NAME>
                </P-PORT-PROTOTYPE>
            </PORTS>
        </ADAPTIVE-APPLICATION-SW-COMPONENT-TYPE>
    </ELEMENTS>
</AR-PACKAGE>

<AR-PACKAGE>
    <SHORT-NAME>Processes</SHORT-NAME>
    <ELEMENTS>
        <PROCESS>
            <SHORT-NAME>fusion_instancel</SHORT-NAME>
        </PROCESS>
    </ELEMENTS>
</AR-PACKAGE>

```

```

    </ELEMENTS>
  </AR-PACKAGE>
</AR-PACKAGE>

  <SHORT-NAME>ServiceInterfaceDeployments</SHORT-NAME>
  <ELEMENTS>
    <SOMEIP-SERVICE-INTERFACE-DEPLOYMENT>
      <SHORT-NAME>radar_Someip</SHORT-NAME>
    </SOMEIP-SERVICE-INTERFACE-DEPLOYMENT>
    <DDS-SERVICE-INTERFACE-DEPLOYMENT>
      <SHORT-NAME>radar_Dds</SHORT-NAME>
    </DDS-SERVICE-INTERFACE-DEPLOYMENT>
  </ELEMENTS>
</AR-PACKAGE>

<AR-PACKAGE>
  <SHORT-NAME>system</SHORT-NAME>
  <ELEMENTS>
    <ETHERNET-CLUSTER>
      <SHORT-NAME>theEthCluster</SHORT-NAME>
      <ETHERNET-CLUSTER-VARIANTS>
        <ETHERNET-CLUSTER-CONDITIONAL>
          <PHYSICAL-CHANNELS>
            <ETHERNET-PHYSICAL-CHANNEL>
              <SHORT-NAME>theEthPhysChannell1</SHORT-NAME>
              <NETWORK-ENDPOINTS>
                <NETWORK-ENDPOINT>
                  <SHORT-NAME>publisherNetworkEndpoint</SHORT-NAME>
                </NETWORK-ENDPOINT>
                <NETWORK-ENDPOINT>
                  <SHORT-NAME>SDNetworkEndpoint</SHORT-NAME>
                </NETWORK-ENDPOINT>
              </NETWORK-ENDPOINTS>
            </ETHERNET-PHYSICAL-CHANNEL>
          </PHYSICAL-CHANNELS>
        </ETHERNET-CLUSTER-CONDITIONAL>
      </ETHERNET-CLUSTER-VARIANTS>
    </ETHERNET-CLUSTER>
  </ELEMENTS>
</AR-PACKAGE>

</AR-PACKAGES>
</AUTOSAR>

```

Listing 6.11: Service Instance Manifest specification of ProvidedServiceInstance

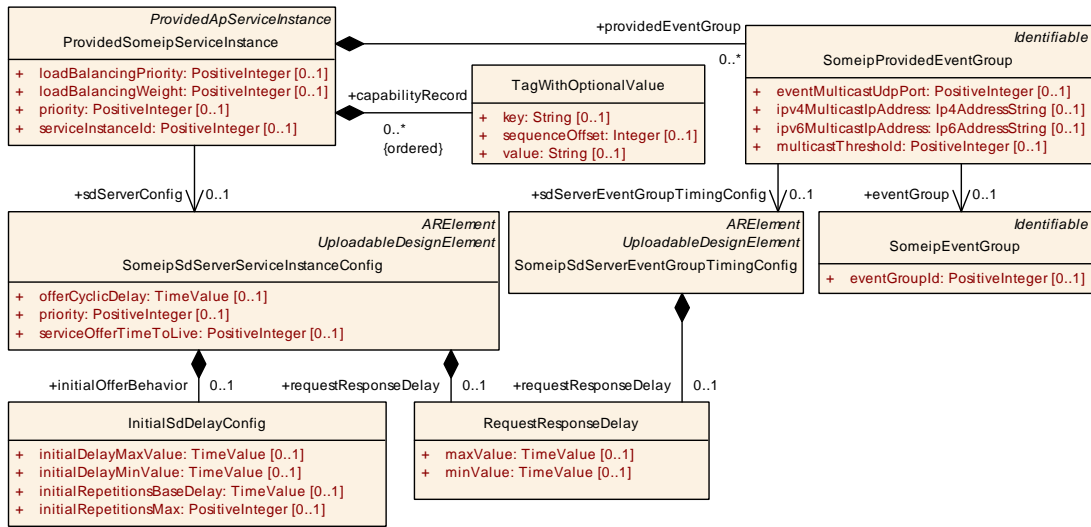


Figure 6.8: SOME/IP Provided Service Instance Deployment

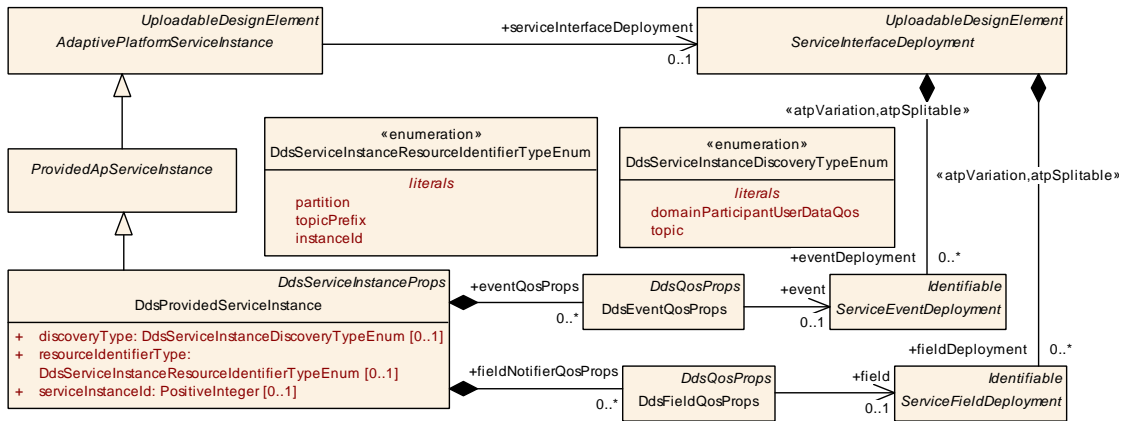


Figure 6.9: DDS Provided Service Instance Deployment

Adaptive Applications looking to bind against instances of the /apd/da/radar SI will now have to properly configure ServiceInstanceToMachineMapping, ServiceInstanceToPortPrototypeMapping and ProvidedServiceInstance subclasses in their Instance Manifests, according to their Network Binding (DDS or SOME/IP) preference.

6.4 Usage of CommunicationGroups

NOTE: Communication Groups have been set to obsolete and removed from [3], therefore it is no longer supported in future AUTOSAR release(s). This section has been left merely as reference.

6.4.1 Setup

This section describes the configuration steps to define a Communication Group by using the Communication Group Template (category = COMMUNICATION_GROUP). There are three items which need to be specified to define a Communication Group:

- The name of the Communication Group = SHORT-NAME.
- The data type of messages sent to the clients.
- The data type of response message sent by the clients.

The figure below presents the Communication Group Template.

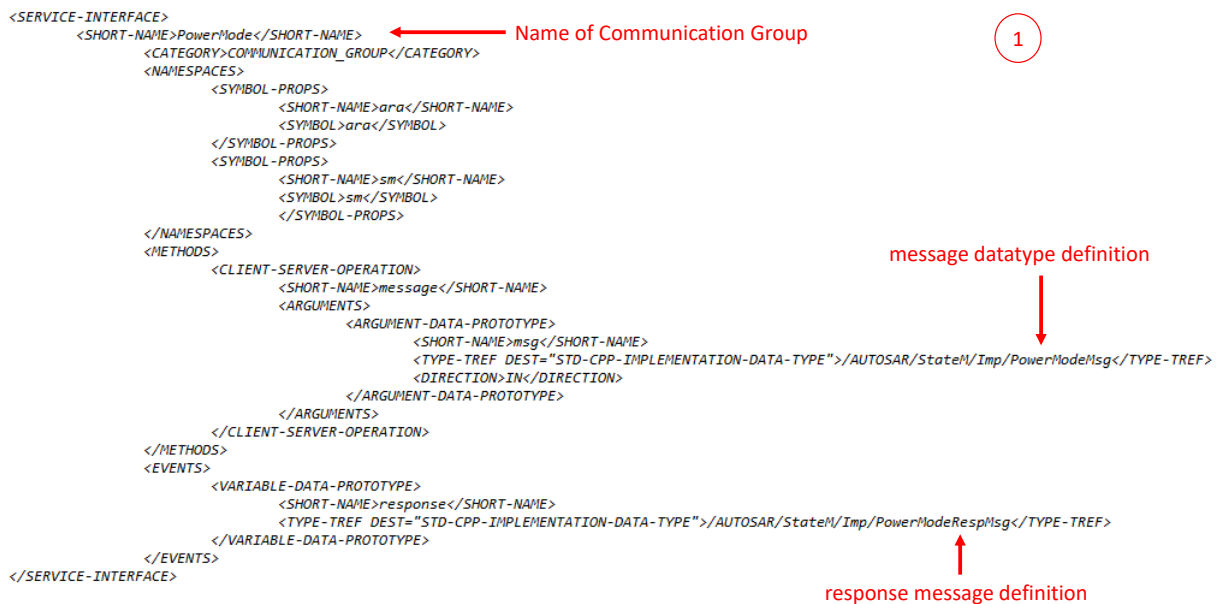


Figure 6.10: CG Template Service Interface Description, example PowerMode

6.4.1.1 Service Interface Description Generation

To create Communication Group server and client ServiceInterface description files, the Communication Group Template need to be setup (see chapter before). Based on the Communication Group Template the ServiceInterface description files of the Communication Group server and client can be generated. See figures below.

Note: The Communication Group Template is a ServiceInterface which is only used for generation of Communication Group ServiceInterfaces and it has else no usage.

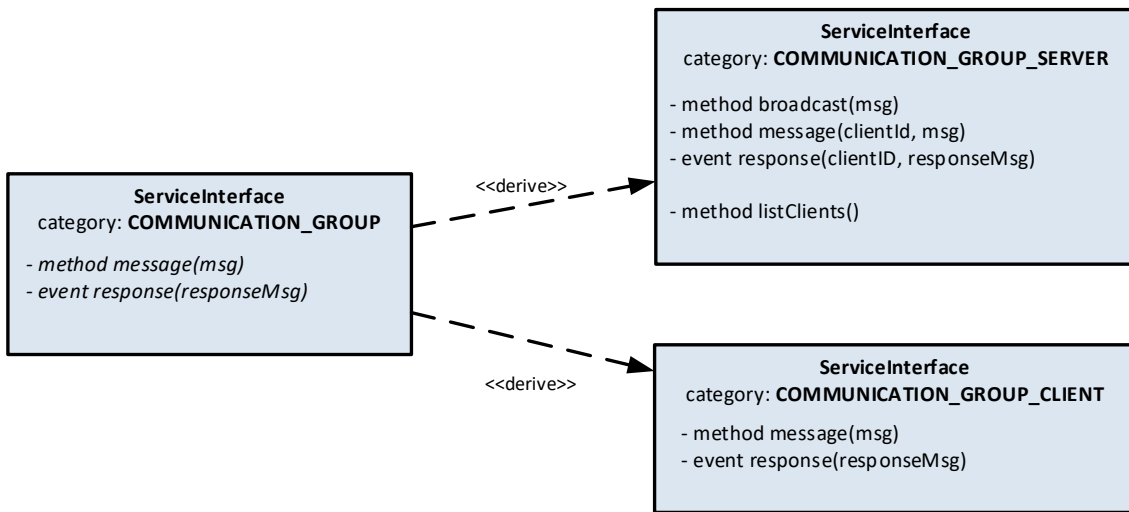


Figure 6.11: Communication Group Service Interfaces Relationship

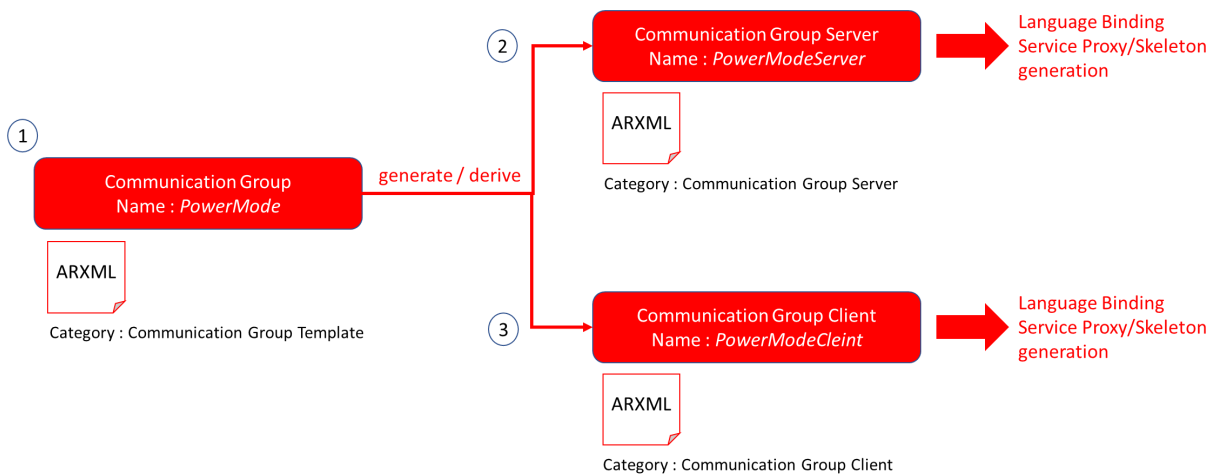


Figure 6.12: Communication Group Service Interfaces Generation

6.4.1.2 Proxy and Skeleton Generation

Using the `ServiceInterface` descriptions for the Communication Group server and client the according Proxy and Skeleton Service classes can be generated (hint: standard `ara::com` [SI](#) generation flow). See figures below.

```

<SERVICE-INTERFACE>
  <SHORT-NAME>PowerModeServer</SHORT-NAME>
  <CATEGORY>COMMUNICATION_GROUP_SERVER</CATEGORY>
  <NAMESPACES>
    <SYMBOL-PROPS>
      <SHORT-NAME>ara</SHORT-NAME>
      <SYMBOL>ara</SYMBOL>
    </SYMBOL-PROPS>
    <SYMBOL-PROPS>
      <SHORT-NAME>sm</SHORT-NAME>
      <SYMBOL>sm</SYMBOL>
    </SYMBOL-PROPS>
  </NAMESPACES>
  <METHODS>
    <CLIENT-SERVER-OPERATION>
      <SHORT-NAME>broadcast</SHORT-NAME>
      <ARGUMENTS>
        <ARGUMENT-DATA-PROTOTYPE>
          <SHORT-NAME>msg</SHORT-NAME>
          <TYPE-TREF DEST="STD-CPP-IMPLEMENTATION-DATA-TYPE"/>/AUTOSAR/StateM/Imp/PowerModeMsg</TYPE-TREF>
          <DIRECTION>IN</DIRECTION>
        </ARGUMENT-DATA-PROTOTYPE>
      </ARGUMENTS>
    </CLIENT-SERVER-OPERATION>
    <CLIENT-SERVER-OPERATION>
      <SHORT-NAME>message</SHORT-NAME>
      <ARGUMENTS>
        <ARGUMENT-DATA-PROTOTYPE>
          <SHORT-NAME>clientID</SHORT-NAME>
          <TYPE-TREF DEST="STD-CPP-IMPLEMENTATION-DATA-TYPE"/>/AUTOSAR/StdTypes/uint32_t</TYPE-TREF>
          <DIRECTION>IN</DIRECTION>
        </ARGUMENT-DATA-PROTOTYPE>
      </ARGUMENTS>
        <ARGUMENTS>
          <ARGUMENT-DATA-PROTOTYPE>
            <SHORT-NAME>msg</SHORT-NAME>
            <TYPE-TREF DEST="STD-CPP-IMPLEMENTATION-DATA-TYPE"/>/AUTOSAR/StateM/Imp/PowerModeMsg</TYPE-TREF>
            <DIRECTION>IN</DIRECTION>
          </ARGUMENT-DATA-PROTOTYPE>
        </ARGUMENTS>
      </CLIENT-SERVER-OPERATION>
    <CLIENT-SERVER-OPERATION>
      <SHORT-NAME>ListClients</SHORT-NAME>
      <ARGUMENTS>
        <ARGUMENT-DATA-PROTOTYPE>
          <SHORT-NAME>clients</SHORT-NAME>
          <TYPE-TREF DEST="STD-CPP-IMPLEMENTATION-DATA-TYPE"/>/AUTOSAR/StateM/Imp/Clients</TYPE-TREF>
          <DIRECTION>IN</DIRECTION>
        </ARGUMENT-DATA-PROTOTYPE>
      </ARGUMENTS>
    </CLIENT-SERVER-OPERATION>
  </METHODS>
  <EVENTS>
    <VARIABLE-DATA-PROTOTYPE>
      <SHORT-NAME>response</SHORT-NAME>
      <TYPE-TREF DEST="STD-CPP-IMPLEMENTATION-DATA-TYPE"/>/AUTOSAR/StateM/Imp/PowerModeResponse</TYPE-TREF>
    </VARIABLE-DATA-PROTOTYPE>
  </EVENTS>
</SERVICE-INTERFACE>

```

2

Figure 6.13: CG Server Service Interface Description, example PowerMode

```

<SERVICE-INTERFACE>
  <SHORT-NAME>PowerModeClient</SHORT-NAME>
  <CATEGORY>COMMUNICATION_GROUP_CLIENT</CATEGORY>
  <NAMESPACES>
    <SYMBOL-PROPS>
      <SHORT-NAME>ara</SHORT-NAME>
      <SYMBOL>ara</SYMBOL>
    </SYMBOL-PROPS>
    <SYMBOL-PROPS>
      <SHORT-NAME>sm</SHORT-NAME>
      <SYMBOL>sm</SYMBOL>
    </SYMBOL-PROPS>
  </NAMESPACES>
  <METHODS>
    <CLIENT-SERVER-OPERATION>
      <SHORT-NAME>message</SHORT-NAME>
      <ARGUMENTS>
        <ARGUMENT-DATA-PROTOTYPE>
          <SHORT-NAME>msg</SHORT-NAME>
          <TYPE-TREF DEST="STD-CPP-IMPLEMENTATION-DATA-TYPE"/>AUTOSAR/StateN/Imp/PowerModeMsg</TYPE-TREF>
          <DIRECTION>IN</DIRECTION>
        </ARGUMENT-DATA-PROTOTYPE>
      </ARGUMENTS>
    </CLIENT-SERVER-OPERATION>
  </METHODS>
  <EVENTS>
    <VARIABLE-DATA-PROTOTYPE>
      <SHORT-NAME>response</SHORT-NAME>
      <TYPE-TREF DEST="STD-CPP-IMPLEMENTATION-DATA-TYPE"/>AUTOSAR/StateN/Imp/PowerModeRespMsg</TYPE-TREF>
    </VARIABLE-DATA-PROTOTYPE>
  </EVENTS>
</SERVICE-INTERFACE>

```

Figure 6.14: CG Client Service Interface Description, example PowerMode

6.4.2 Example

The following chapter outlines the steps to define and create the Communication Group [PowerMode](#).

- Define the CG name and add this name as `SHORT-NAME` into the `ServiceInterface` description of the Communication Group Template.

+ `SHORT-NAME` = `PowerMode`

- Define the message data type for peer to peer and broadcast message and edit it into the `ServiceInterface` description of the Communication Group Template.

+ message definition `<typename T>`

```
PowerModeMsg = [On, Off, Suspend]
```

background info:

sample CG code:

```
template <typename T>
ara::core::Future<void> broadcast (const T& msg);
```

sample CG code:

```
template <typename T>
ara::core::Future<void> message (uint32_t clientID, const T& msg);
```

sample CG code:

```
template <typename T>
ara::core::Future<void> message (const T& msg);
```

- Define the response message data type and edit it into the ServiceInterface description of the Communication Group Template.

+ response message definition <typename R>

```
PowerModeRspMsg = [Done, Failed, Busy, NotSupported]
```

background info:

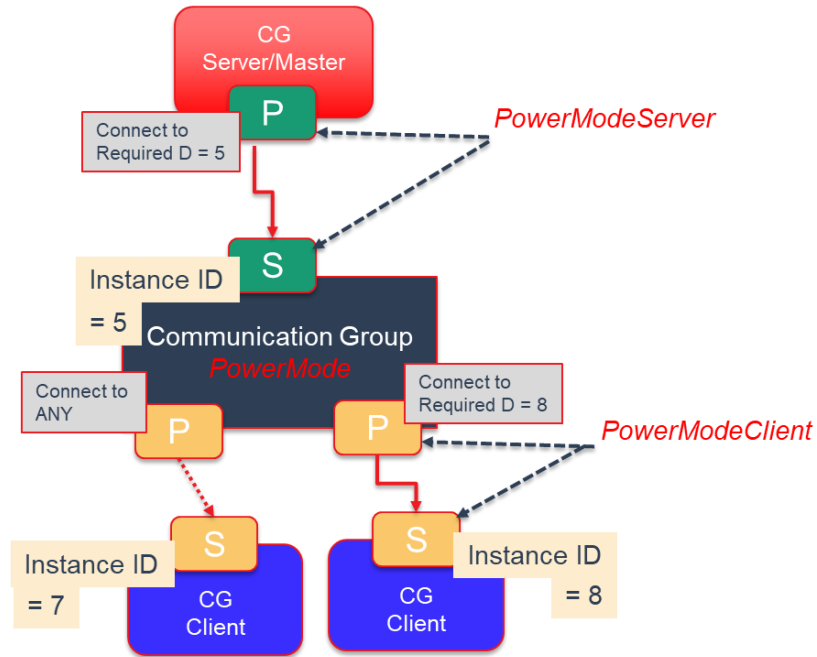
sample CG code:

```
template <typename R>
struct Response {
    uint32_t      clientID;
    const R&      responseMsg }
```

sample CG code:

```
template <typename R>
const R& responseMsg;
```

- Generate the Server and Client ServiceInterface descriptions out of the setup Communication Group Template.
- Generate the Server and Client ServiceInterface Proxy and Skeleton Service class templates.
- Instantiate the Communication Group Server and Client Proxy and Skeleton Service classes into the application code. See examples below.



Confidential C

Figure 6.15: Usage Example Communication Group PowerMode

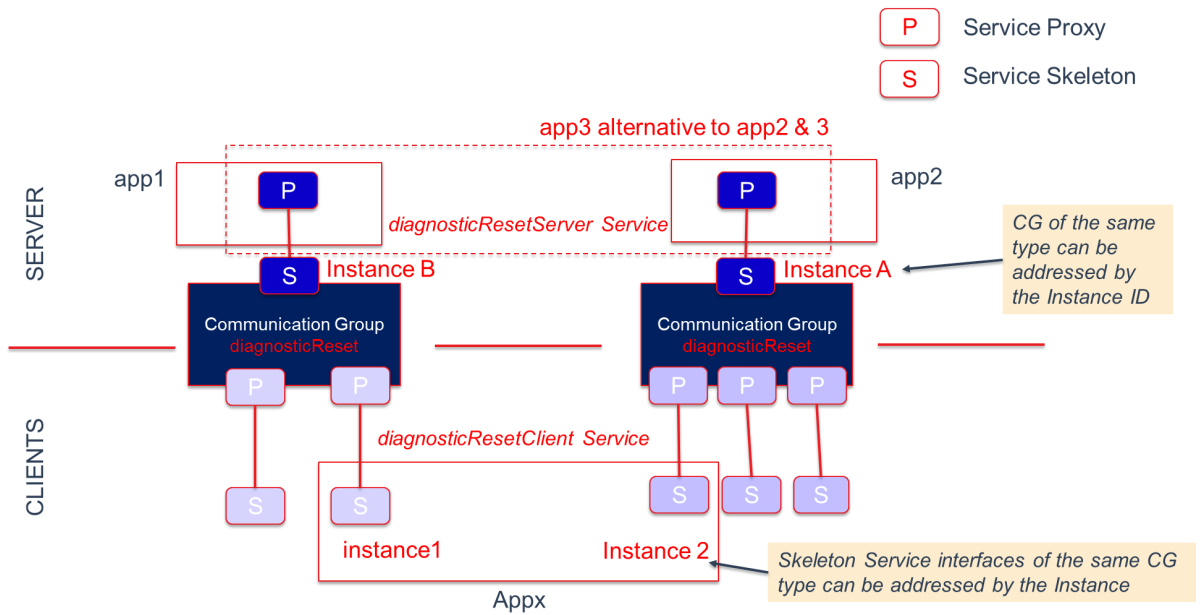


Figure 6.16: Usage Example 2 x Communication Groups diagnosticReset

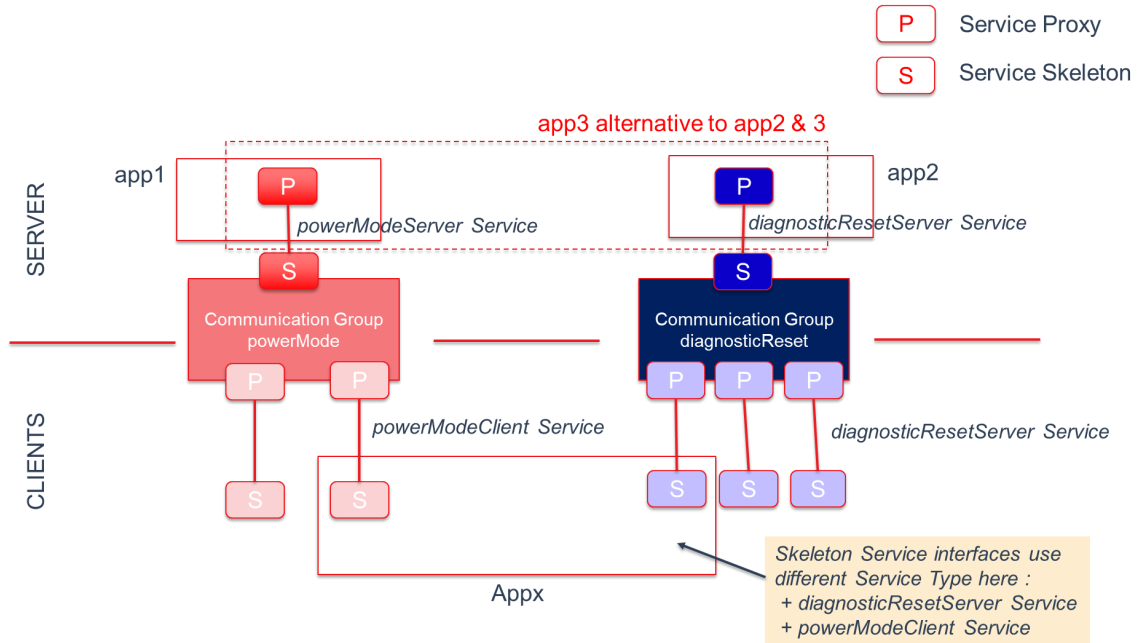


Figure 6.17: Usage Example Communication Group PowerMode and diagnosticReset

7 Appendix

7.1 Serialization

`Serialization` (see [11]) is the process of transforming certain data structures into a standardized format for exchange between a sender and a (possibly different) receiver. You typically have this notion if you transfer data from one network node to another. When putting data on the wire and reading it back, you have to follow exact, agreed-on rules to be able to correctly interpret the data on the receiver side. For the network communication use case the need for a defined approach to convert an in-process data representation into a wire-format and back is very obvious: The boxes doing the communication might be based on different micro-controllers with different endianness and different data-word sizes (16-bit, 32-bit, 64-bit) and therefore employing totally different alignments. In the AUTOSAR CP `serialization` typically plays no role for platform internal/node internal communication! Here the internal in-memory data representation can be directly copied from a sender to a receiver. This is possible, because three assumptions are made in the typical CP product:

- Endianness is identical among all local SWCs.
- Alignment of certain data structures is homogeneous among all local SWCs.
- Data structures exchanged are contiguous in memory.

The first point is maybe a bit pathological as it is most common, that “internal” communication generally means communication on a single- or multi-core MCU or even a multi-processor system, where endianness is identical everywhere. Only if we look at a system/node composed of CPUs made of different micro-controller families this assumption may be invalid, but then you are already in the discussion, whether this communication is still “internal” in the typical sense. The second assumption is valid/acceptable for CP as here a static image for the entire single address space system is built from sources and/or object files, which demands that compiler settings among the different parts of the image are somewhat aligned anyway. The third one is also assured in CP. It is not allowed/possible to model non contiguous data types, which get used in inter-SWC communication.

For the AP things look indeed different. Here the loading of executables during runtime, which have been built independently at different times and have been uploaded to an AP ECU at different times, is definitely a supported use case. The chance, that compiler settings for different `ara::com` applications were different regarding alignment decisions is consequently high. Therefore an AP product (more concrete its `IPC` binding implementation) has to use/support serialization of exchanged event/field-/method data. How serialization for AP internal `IPC` is done (i.e. to what generalized format) is fully up to the AP vendor. Also regarding the 3rd point, the AP is less restrictive. So for example the AP supports exchange of `std::map` data types or record like datatypes, which contain variable-length members. These datatypes are generally NOT contiguous in-memory (depending on the allocation strategy). So even if the data contained in the map or records is compatible with the receiver layout wise, a deep

copy (meaning collecting contained elements and their references from various memory regions — see [12]) must be done during transfer. Of course the product vendor could apply optimization strategies to get rid of the serialization and de-serialization stages within a communication path:

- Regarding alignment issues, the most simple one could be to allow the integrator of the system to configure, that alignment for certain communication relations can be considered compatible (because he has the needed knowledge about the involved components).
- Another approach common to middleware technology is to verify, whether alignment settings on both sides are equal by exchanging a check-pattern as kind of an init-sequence before first `ara::com` communication call.
- The problem regarding need for deep-copying because of non-contiguous memory allocation could be circumvented by providing vector implementations which care for continuity.

7.1.1 Zero-Copy implications

One thing which typically is at the top of the list of performance optimizations in IPC/middleware implementations is the avoidance of unnecessary copies between sender and the receiver of data. So the buzzword “zero-copy” is widely used to describe this pattern. When we talk about AP, where we have architectural expectations like applications running in separate processes providing memory protection, the typical communication approach needs at least ONE copy of the data from source address space to target address space. Highly optimizing middleware/IPC implementations could even get rid of this single copy step by setting up shared memory regions between communicating `ara::com` components. If you look at [Listing 5.19](#), you see, that we directly encourage such implementation approaches in the API design. But the not so good news is, that if the product vendor does NOT solve the serialization problem, he barely gets benefit from the shared memory approach: If conversions (aka de/serialization) have to be done between communication partners, copying must be done anyhow — so tricky shared memory approaches to aim for “zero-copy” do not pay.

7.2 Service Discovery Implementation Strategies

As laid out in the preceding chapters, `ara::com` expects the functionality of a service discovery being implemented by the product vendor. As the service discovery functionality is basically defined at the API level with the methods for `FindService`, `OfferService` and `StopOfferService`, the protocol and implementation details are partially open.

When an AP node (more concretely an AP SWC) offers a service over the network or requires a service from another network node, then service discovery/service registry obviously takes place over the wire. The protocol for service discovery over the wire needs to be completely specified by the used communication protocol. For SOME/IP, this is done in the SOME/IP Service Discovery Protocol Specification [13]. But if an `ara::com` application wants to communicate with another `ara::com` application on the same node within the AP of the same vendor there has to be a local variant of a service discovery available. Here the only difference is, that the protocol implementation for service discovery taking place locally is totally up to the AP product vendor.

7.2.1 Central vs Distributed approach

From an abstract perspective a AP product vendor could choose between two approaches: The first one is a centralist approach, where the vendor decides to have one central entity (f.i. a daemon process), which:

- maintains a registry of all service instances together with their location information
- serves all `FindService`, `OfferService` and `StopOfferService` requests from local `ara::com` applications, thereby either updating the registry (`OfferService`, `StopOfferService`) or querying the registry (`FindService`)
- serves all SOME/IP SD messages from the network either updating its registry (`SOME/IP Offer Service received`) or querying the registry (`SOME/IP Find Service received`)
- propagates local updates to its registry to the network by sending out SOME/IP SD messages.

The following figure roughly sketches this approach.

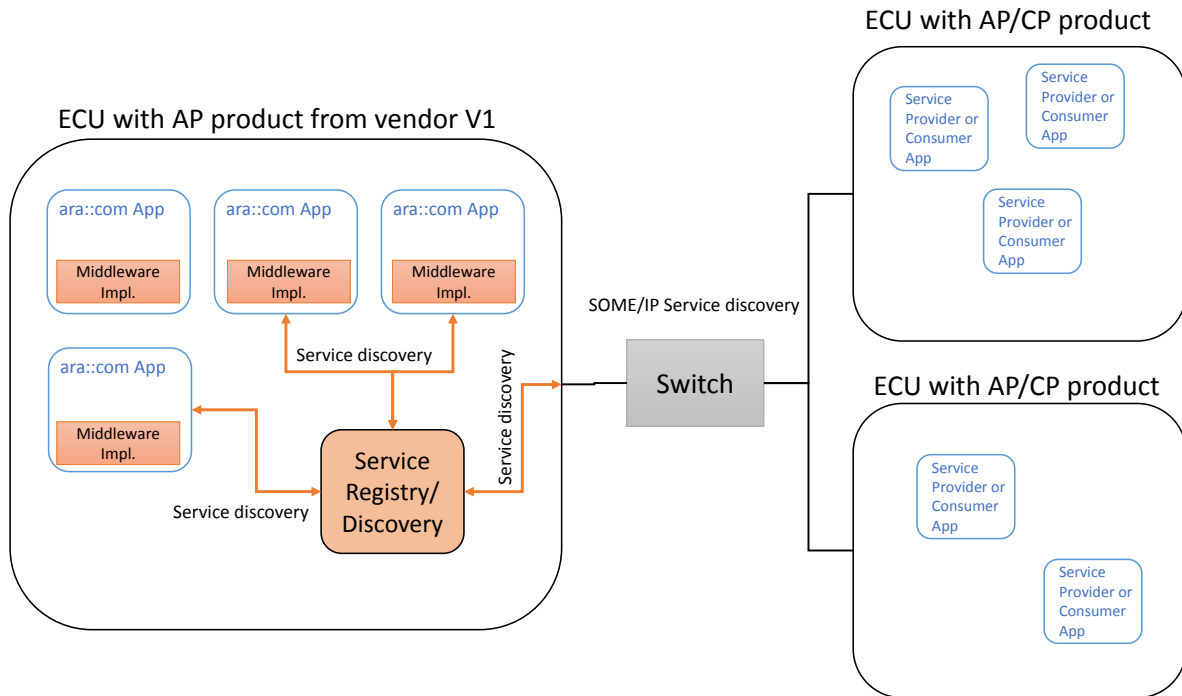


Figure 7.1: Centralized discovery approach

A slightly different — more distributed — approach would be, to distribute the service registry information (availability and location information) among the `ara::com` applications within the node. So for the node local communication use case no prominent discovery demon would be needed. That could be technically reached by having a broadcast-like communication. That means any service offering and finding is propagated to all local `ara::com` applications, so that each application has a local (in process) view of the service registry. There might be a benefit with this approach as local communication might be more flexible/stable as it is not dependent from a single registry demon. However, for the service discovery communication to/from the network a single responsible instance is needed anyhow. Here the distributed approach is not feasible as SOME/IP SD requires a fixed/defined set of ports, which just can be provided (in typical operating systems / with typical network stacks) by a single application process.

At the end we also do have a singleton/central instance, with the slight difference, that it is responsible for taking the role as a service discovery protocol bridge between node local discovery protocol and network SOME/IP SD protocol. On top of that — since registry is duplicated/distributed among all `ara::com` applications within the node — this bridge also holds a local registry.

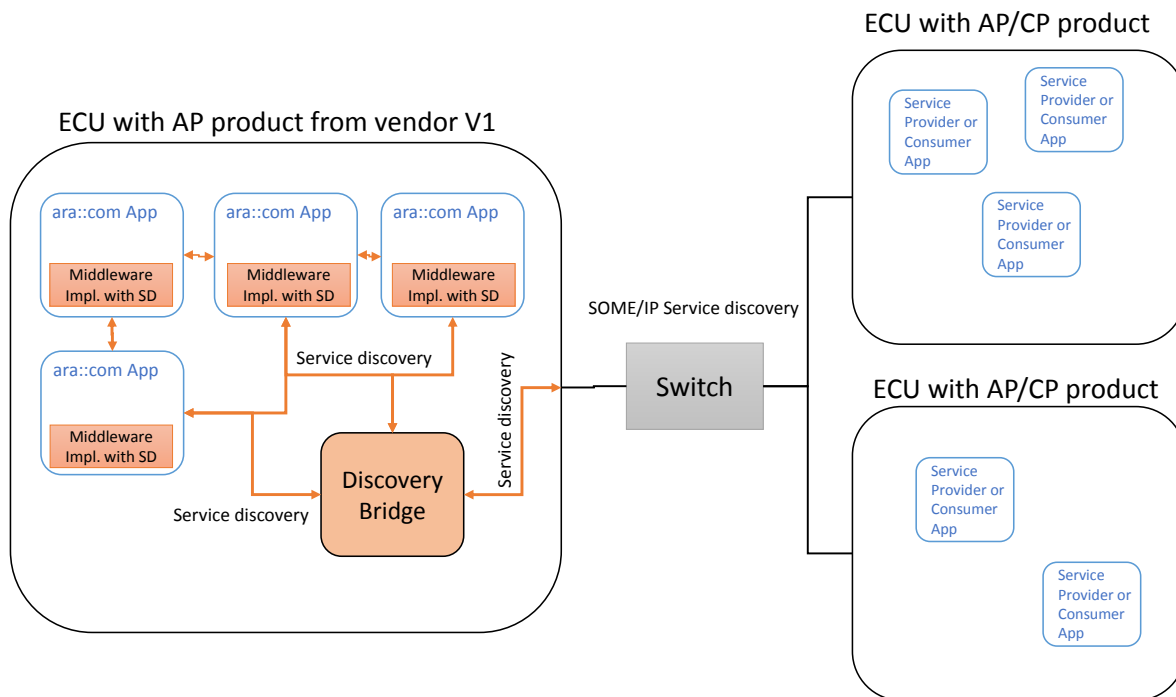


Figure 7.2: Distributed discovery approach

7.3 Multi-Binding implications

As shortly discussed in [Section 5.4.3 Multi-Binding](#) describes the solution to support setups, where the technical transport/connection between different instances of a certain proxy class/skeleton class are different. There might be various technical reasons for that:

- proxy class uses different transport/IPC to communicate with different skeleton instances. Reason: Different service instances support different transport mechanisms because of deployment decisions.
- symmetrically it may also be the case, that different proxy instances for the same skeleton instance uses different transport/IPC to communicate with this instance: The skeleton instance supports multiple transport mechanisms to get contacted.

7.3.1 Simple Multi-Binding use case

The following figure depicts an obvious and/or rather simple case. In this example, which only deals with node local (inside one AP product/ECU) communication between service consumers (proxy) and service providers (skeleton), there are two instances of the same proxy class on the service consumer side. You see in the picture, that the service consumer application has triggered a “FindService” first, which returned two handles for two different service instances of the searched service type. The service consumer application has instantiated a proxy instance for each of those handles. Now in this example the instance 1 of the service is located inside the same adaptive application (same process/address space) as the service consumer (proxy instance 1), while the service instance 2 is located in a different adaptive application (different process/address space).

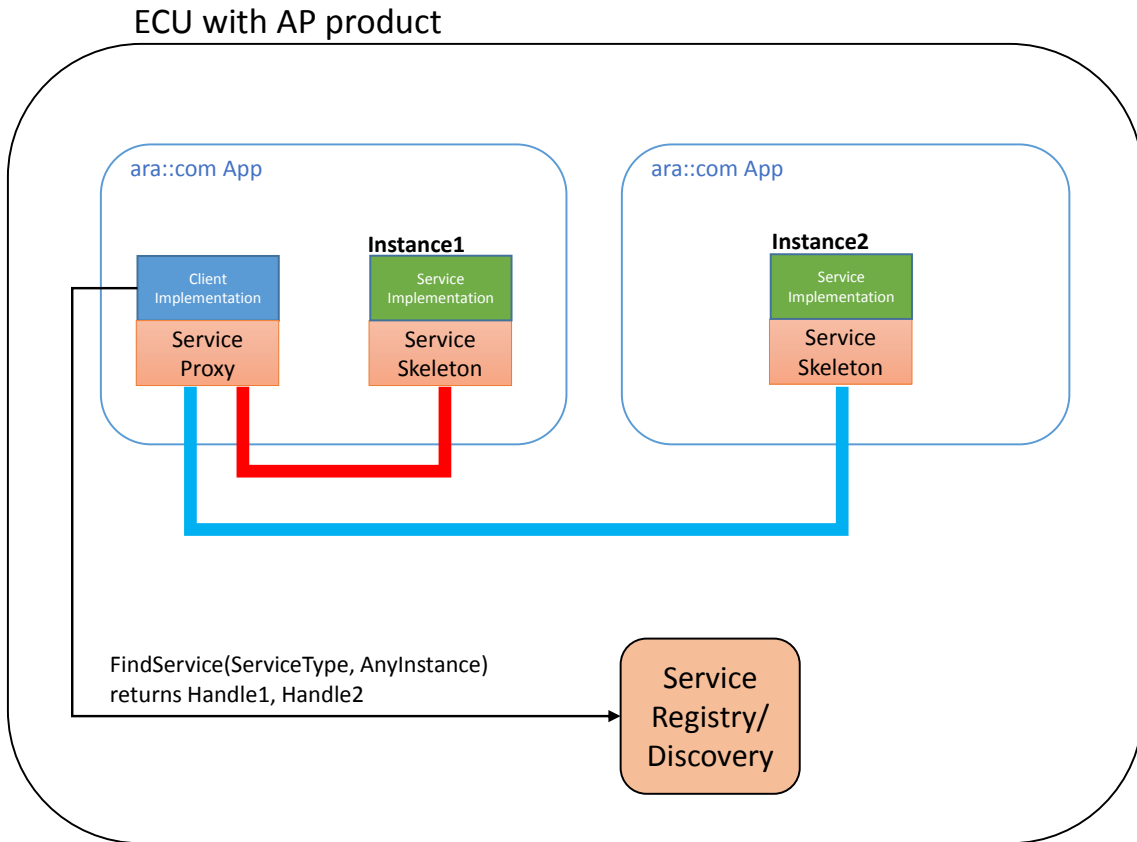


Figure 7.3: Simple Multi-Binding intra AP example

The line symbolizing the transport layer between proxies and skeletons are colored differently in this picture: The instance of the proxy class for instance 1 has a red colored transport layer (binding implementation), while the transport layer for instance 2 is colored blue. They are colored differently because the used technology will be different already on the level of the proxy implementation. At least if you expect that the AP product vendor (in the role as `IPC` binding implementer) strives for a well performing product!

The communication between proxy instance 1 and the service instance 1 (red) should in this case be optimized to a plain method call, since proxy instance and skeleton instance 1 are contained in ONE process.

The communication between proxy instance 2 and the service instance 2 (blue) is a real IPC. So the actions taken here are of much higher costs involving most likely a variety of syscalls/kernel context switches to transfer calls/data from process of service consumer application to service application (typically using basic technologies like pipes, sockets or shared mem with some signaling on top for control).

So from the service consumer side application developer it is totally transparent: From the vendors `ProxyClass::FindService` implementation he gets two opaque handles for the two service instances, from which he creates two instances of the same

proxy class. But “by magic” both proxies behave totally different in the way, they contact their respective service instances. So — somehow there must be some information contained inside this handle, from which the proxy class instance knows which technical transport to choose. Although this use case looks simple at the first look it isn’t on the second ... The question is: *Who* writes *When* into the handle, that the proxy instance created from it shall use a direct method/function call instead of a more complex IPC mechanism or vice versa?

At the point in time when instance 1 of the service does register itself via `SkeletonClass::OfferService` at the registry/service discovery, this cannot be decided! Since it depends on the service consumer which uses it later on. So most likely the `SkeletonClass::OfferService` implementation of the AP vendor takes the needed information from the argument (skeleton generated by the AP vendor) and notifies via AP vendor specific IPC the registry/service discovery implementation of the AP vendor. The many “AP vendor” in the preceding sentence were intentional. Just showing, that all those mechanisms going on here are not standardized and can therefore deliberately designed and optimized by the AP vendors. However, the basic steps will remain. So what typically will be communicated from the service instance side to the registry/discovery in the course of `SkeletonClass::OfferService` is the technical addressing information, how the instance could be reached via the AP products local IPC implementation.

Normally there will be only ONE IPC-mechanism used inside one AP product/AP node! If the product vendor already has implemented a highly optimized/efficient local IPC implementation between adaptive applications, which will then be generally used. So — in our example let’s say the underlying IPC-mechanism is unix domain sockets — the skeleton instance 1 would get/create some file descriptor to which its socket endpoint is connected and would communicate this descriptor to the registry/service discovery during `SkeletonClass::OfferService`. Same goes for the skeleton instance 2, just the descriptor is different. When later on the service consumer application part does a `ProxyClass::FindService`, the registry will send the addressing information for both service instances to the service consumer, where they are visible as two opaque handles.

So in this example obviously the handles look exactly the same — with the small difference, that the contained file descriptor values would be different as they reference distinctive unix domain sockets. So in this case it somehow has to be detected inside the proxy for instance 1, that there is the possibility to optimize for direct method/function calls. One possible trivial trick could be, that inside the addressing information, which skeleton instance 1 gives to the registry/discovery, also the ID of the process (pid) is contained; either explicitly or by including it into the socket descriptor filename. So the service consumer side proxy instance 1 could simply check, whether the PID inside the handle denotes the same process as itself and could then use the optimized path. By the way: Detection of process local optimization potential is a triviality, which almost every existing middleware implementation does today — so no further need to stress this topic.

Now, if we step back, we have to realize, that our simple example here does NOT fully reflect what `Multi-Binding` means. It does indeed describe the case, where two instances of the same proxy class use different transport layers to contact the service instance, but as the example shows, this is NOT reflected in the handles denoting the different instances, but is simply an optimization! In our concrete example, the service consumer using the proxy instance 1 to communicate with the service instance 1 could have used also the Unix domain socket transport like the proxy instance 2 without any functional losings — only from a non-functional performance viewpoint it would be obviously bad. Nonetheless this simple scenario was worth being mentioned here as it is a real-world scenario, which is very likely to happen in many deployments and therefore must be well supported!

7.3.2 Local/Network Multi-Binding use case

After we have seen a special variant of `Multi-Binding` in the preceding section, we now look at a variant, which can also be considered as being a real-world case. Let's suppose, we have have a setup quite similar to the one of the preceding chapter. The only difference is now, that the instance 2 of the service is located on a different ECU attached to the same Ethernet network as our ECU with the AP product, where the service consumer (with its proxies for instance 1 and 2) resides. As the standard protocol on Ethernet for AP is SOME/IP, it is expected, that the communication between both ECUs is based on SOME/IP. For our concrete example this means, that proxy 1 talks to service 1 via unix domain sockets (which might be optimized for process local communication to direct method calls, if the AP vendor/IPC implementer did his homework), while the proxy 2 talks to service 2 via network sockets in a SOME/IP compliant message format.

Before someone notes, that this is not true for the typical SOME/IP deployment, because there adaptive SWCs will not directly open network socket connections to remote nodes: We will cover this in more detail here ([Section 7.3.3](#)), but for now suppose, that this is a realistic scenario. (For other network protocols it might indeed be realistic)

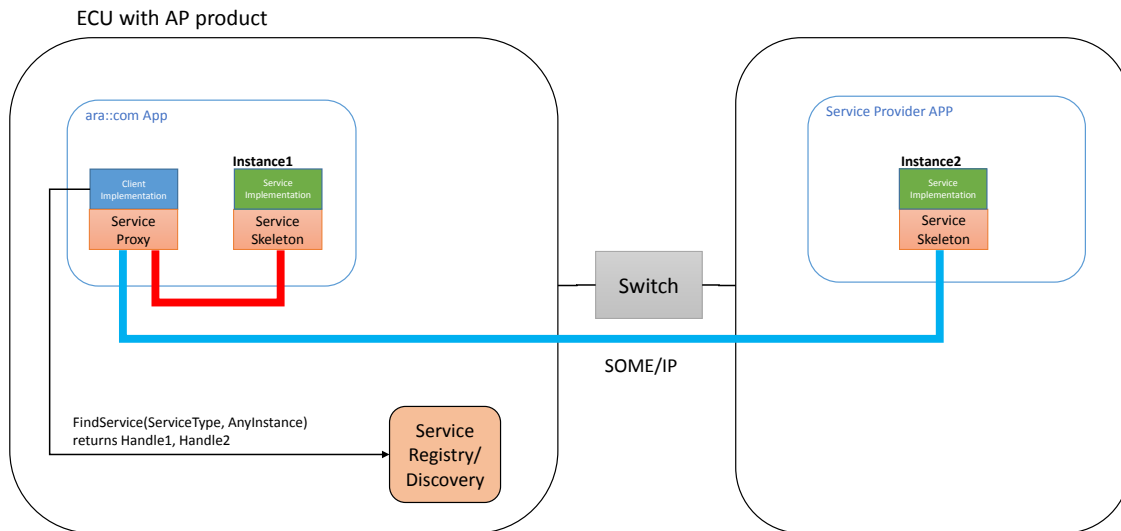


Figure 7.4: Multi-Binding local and network example

So in this scenario the registry/service discovery demon on our AP ECU has seen a service offer of instance 2 and this offer contained the addressing information on IP network endpoint basis. Regarding the service offer of the instance 1 nothing changed: This offer is still connected with some Unix domain socket name, which is essentially a filename. In this example the two handles for instance 1 and 2 returned from `ProxyClass::FindService` internally look very different: Handle of instance 1 contains the information, that it is a Unix domain socket and a name, while handle 2 contains the information, that it is a network socket and an IP address and port number. So — in contrast to our first example ([Section 7.3.1](#)) here we do really have a full blown `Multi-Binding`, where our proxy class `ctor` instantiates/creates two completely different transport mechanisms from handle 1 and handle 2! How this dynamic decision, which transport mechanism to use, made during call of the `ctor`, is technically solved is — again — up to the middleware implementer: The generated proxy class implementation could already contain any supported mechanism and the information contained in the handle is just used to switch between different behavior or the needed transport functionality aka binding could be loaded during runtime after a certain need is detected from the given handle via shared library mechanisms.

7.3.3 Typical SOME/IP Multi-Binding use case

In the previous section we briefly mentioned, that in a typical deployment scenario with SOME/IP as network protocol, it is highly unlikely that an adaptive SWC (i.e. the language and network binding which runs in its context) opens socket connections itself to communicate with a remote service. Why is it unlikely? Because SOME/IP was explicitly designed to use as few ports as possible. The reason for that requirement

comes from low power/low resources embedded ECUs: Managing a huge amount of IP sockets in parallel means huge costs in terms of memory (and runtime) resources. So somehow our AUTOSAR CP siblings which will be main communication partner in an inside vehicle network demand this approach, which is uncommon, compared to non-automotive IT usage pattern for ports.

Typically this requirement leads to an architecture, where the entire SOME/IP traffic of an ECU / network endpoint is routed through one IP port! That means SOME/IP messages originating from/dispatched to many different local applications (service providers or service consumers) are (de)multiplexed to/from one socket connection. In Classic AUTOSAR (CP) this is a straight forward concept, since there is already a shared communication stack through which the entire communication flows. The multiplexing of different upper layer PDUs through one socket is core functionality integrated in CPs SoAd basic software module. For a typical POSIX compatible OS with POSIX socket API, multiplexing SOME/IP communication of many applications to/from one port means the introduction of a separate/central (demon) process, which manages the corresponding port. The task of this process is to bridge between SOME/IP network communication and local communication and vice versa.

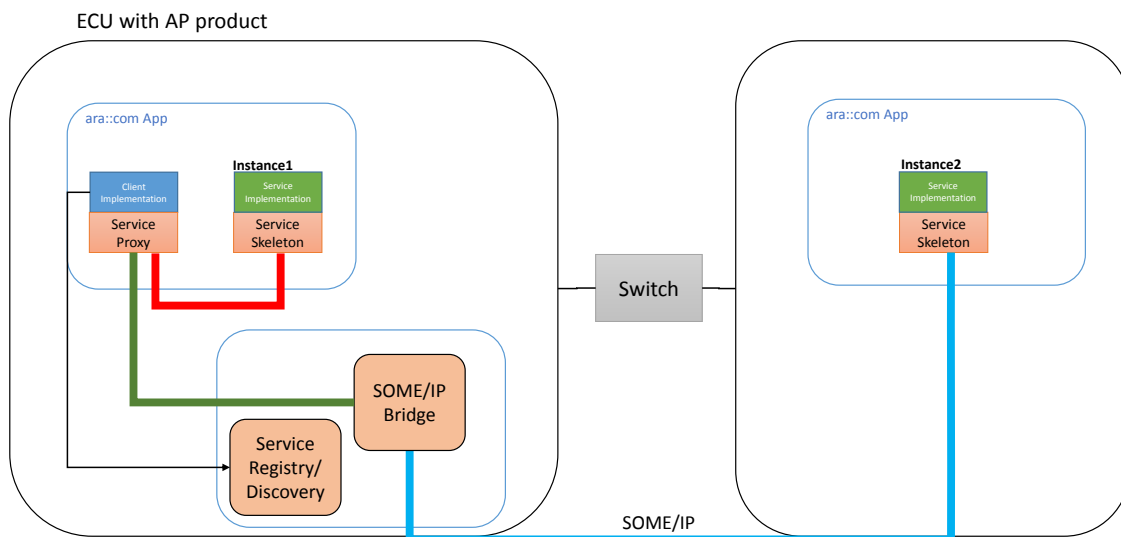


Figure 7.5: SOME/IP Bridge

In the above figure you see, that the service proxy within our `ara::com` enabled application communicates through (green line) a SOME/IP Bridge with the remote service instance 2. Two points which may pop out in this figure:

- we intentionally colored the part of the communication route from app to bridge (green) differently than the part from the bridge to the service instance 2 (blue).
- we intentionally drew a box around the function block service discovery and SOME/IP bridge.

The reason for coloring first part of the route differently from the second one is simple: Both parts use a different transport mechanism. While the first one (green) between the proxy and the bridge uses a fully vendor specific implementation, the second one (blue) has to comply with the SOME/IP specification. “Fully vendor specific” here means, that the vendor not only decides which technology he uses (pipes, sockets, shared mem, ...), but also which serialization format (see [Section 7.1](#)) he employs on that path. Here we obviously dive into the realm of optimizations: In an optimized AP product, the vendor would not apply a different (proprietary) serialization format for the path denoted with the green line. Otherwise it would lead to an inefficient runtime behavior. First the proxy within the service consumer app would employ a proprietary serialization of the data before transferring it to the bridge node and then the bridge would have to de-serialize and re-serialize it to SOME/IP serialization format! So even if the AP product vendor has a much more efficient/refined serialization approach for local communication, using it here does not pay, since then the bridge is not able to simply copy the data through between internal and external side. The result is, that for our example scenario we eventually do have a `Multi-Binding` setup. So even if the technical transport (pipes, unix domain sockets, shared mem, ...) for communication to other local `ara::com` applications and to the bridge node is the same, the serialization part of the binding differs.

Regarding the second noticeable point in the figure: We drew a box around the service discovery and SOME/IP bridge functionality since in product implementations it is very likely, that it is integrated into one component/running within one (demon) process. Both functionalities are highly related: The discovery/registry part also consists of parts local to the ECU (receiving local registrations/offers and serving local Find-Service requests) and network related functions (SOME/IP service discovery based offers/finds) , where the registry has to arbitrate. This arbitration in its core is also a bridging functionality.

7.4 ara::com and AUTOSAR meta-model relationship

Throughout this document we paid attention to explain `ara::com` API ideas and mechanisms **without** relating to the concrete/specific AP meta-model (the manifest parts of it), which is the basis to formally describe the `SI` signature (and partially the behavior) from which the `ara::com` API artifacts like `ProxyClass` and `SkeletonClass` and data types used in the communication are generated/created. In [Listing 5.1](#) we even introduced an oversimplified/synthetic IDL, just to shield the reader from complexities of the real meta-model/IDL, which wouldn't have added any value at that point.

This chapter shall by no means serve as a thorough explanation of the AUTOSAR meta-model, which is fully described in its own document, but it shall shed some light on the relation between `ara::com` and the meta-model parts described in [1]. So bear in mind, that the following parts are still somewhat high level and try to give a basic understanding of the relationship.

7.4.1 Connection to AUTOSAR_TR_AdaptiveMethodology

Overview of Modeling elements and how they are related to each other: [SI](#), Deployment, Actual generation dependant from provided Deployment Information (E.g. also [SI](#) Elements that will be generated later and connection to Service Instance Manifest)

AUTOSAR Adaptive Platform methodology explains the process aspects necessary to build an Adaptive AUTOSAR system and how they relate to each other [TR_AMETH_00100]. It defines activities and work products delivered or consumed [TR_AMETH_00102] and the Roles performed by OEMs and suppliers.

Major steps involved in the development of Adaptive Software are

- Architecture and Design
- Adaptive Software Development
- Integration and Deployment

Adaptive applications run on top of ARA layer and exchanges the information using [SIs](#) and Ports. Important contribution for [ara::com](#) API work performed during the Integration and Deployment step of Adaptive Methodology. It supports the generation of [SI](#) Description ARXML file, which aggregates the [SIs](#) and ports. [SIs](#) for service-oriented communication defined by Events, Methods and Fields [Listing 5.1]. This is done independent of Software components or Transport layer used for underlying communication.

Adaptive Platform supports two types of ports namely Provided and Required. [SI](#) along with Provided port details used for the generation of the Service Skeleton class and Required port details used for the generation of Proxy classes [Figure 5.2]. Proxy and Skeleton classes use [ara::com](#) API to communicate with other Adaptive Platform clusters and Adaptive Applications.

Service instances are configured, notably the binding of the [SIs](#) to a chosen transport layer, whether a specific service instance is either Provided or Required and whether there is a mapping to a dedicated Machine. The configurations of the service instance are manifested in the Service Instance Manifest.

Executable of an Adaptive Software are instantiated by means of the Execution Manifest. Instantiation here means to bind the executables to the context of specific processes of the operating system. Each process may start with a different start-up configuration depending on a machine mode. Further on, the Execution Manifest also defines Software process dependencies.

7.4.2 Service Interface

The most important meta-model element from the [ara::com](#) perspective is the [SI](#). Most important, because it defines everything signaturewise of an [ara::com](#) proxy or skeleton. The [SI](#) describes the methods, fields and the methods a [SI](#) consists of and

how the signatures of those elements (arguments and data types) look like. So the [Listing 5.1](#) is basically a simplification of meta-model [SI](#) and the real meta-model data type system.

The relationship between the meta-model element [SI](#) and [ara::com](#) is therefore clear: [ara::com](#) proxy and skeleton classes get generated from [SI](#).

7.4.3 Software Component

With software components, the AUTOSAR methodology defines a higher order element than just interfaces. The idea of a software component is to describe a reusable part of software with well defined interfaces. For this the AUTOSAR manifest specification defines a model element `SoftwareComponentType`, which is an abstract element with several concrete subtypes, of which the subtype `AdaptiveApplicationSwComponentType` is the most important one for Adaptive Application software developers. A `SoftwareComponentType` model element is realized by C++ code. Which [SIs](#) such a component "provides to" or "requires from" the outside is expressed by `ports`. Ports are typed by [SIs](#). P-ports express that the [SI](#), which types the port, is provided, while R-ports express, that the [SI](#), which types the port, is required by the `SoftwareComponentType`.

The figure [Figure 7.6](#) gives a coarse idea, how the model view relates to the code implementation.

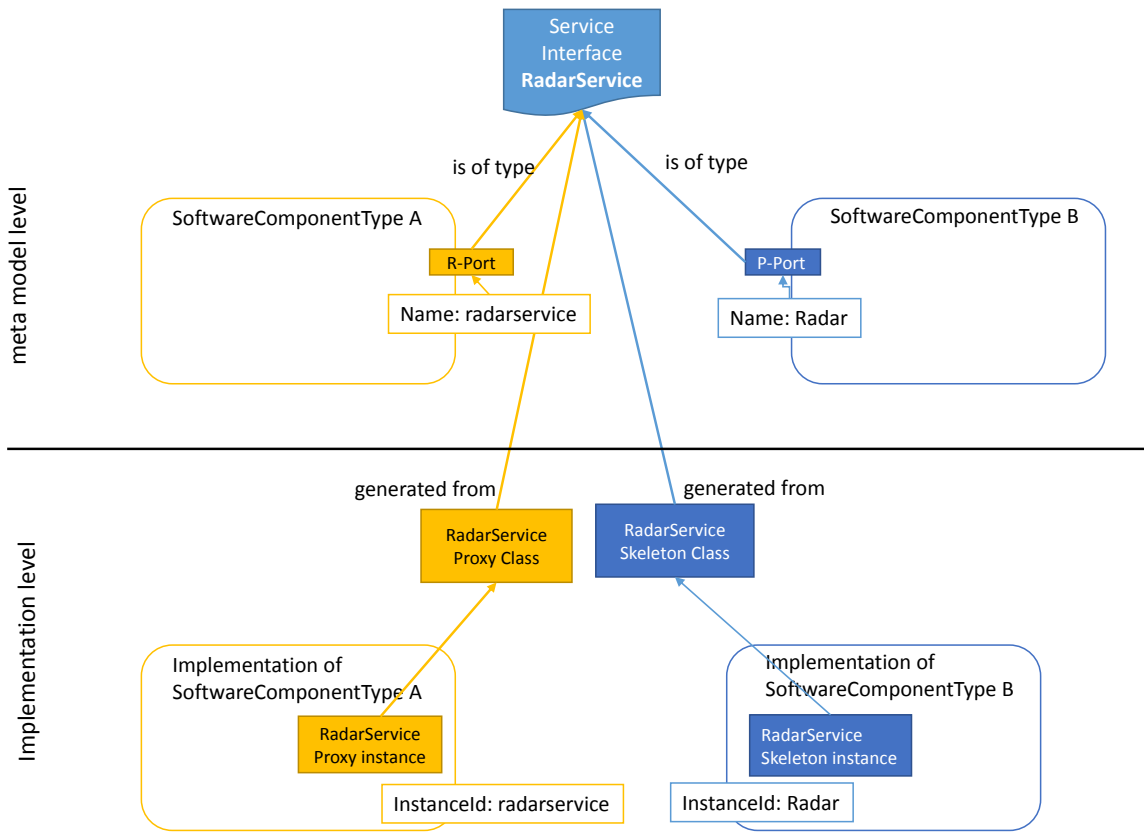


Figure 7.6: meta-model to Implementation

For both of the different `SoftwareComponentTypes` A and B from the example in the upper part (meta-model level) a concrete implementation exists on implementation level (lower part in the figure). The realization/implementation of `R-Port` of `SoftwareComponentType A` is based on an instance of `ara::com` proxy class on implementation level, while the `P-Port` implementation of `SoftwareComponentType B` is using an instance of `ara::com` skeleton class. Proxy and skeleton class are generated from the `SI` definition `SI`, which is referenced by the corresponding ports. In this example it is the `SI` "RadarService", which we already use throughout the document.

Such a code fragment, which realizes a `SoftwareComponentType` can obviously be re-used. On C++ implementation level an implementation of an `AdaptiveApplicationSwComponentType` typically boils down to one or several C++ classes. So re-use simply means instantiating this class/those classes in different contexts multiple times. Here we can basically distinguish the following cases:

- Explicit multiple instantiation of the C++ class(es) within Code.
- Implicit multiple instantiation by starting/running the same executable multiple times.

The first case still belongs to the realm of "implementation level".




Figure 7.7: Multiple Instantiation in Implementation Contexts

The figure above shows an arbitrary example, where the implementations of A and B are instantiated in different contexts. On the lower left side there is an Executable 1, which directly uses two instances of As impl and one instance of Bs impl. Opposed to that, the right side shows an Executable 2, which "directly" (i.e. on its top most level) uses one instance of Bs impl and an instance of a composite software component, which itself "in its body" again instantiates one instance of As and Bs impl. Note: This natural implementation concept of composing software components from other components to a bigger/composite artefact is fully reflected in the AUTOSAR meta-model in the form of a `CompositionSwComponentType`, which itself is a `SoftwareComponentType` and allows arbitrary recursive nesting/compositing of software components.

The second case on the other hand belongs to the realm of "deployment level" and shall be clarified in the following sub-chapter.

7.4.4 Adaptive Application/Executables and Processes

Deployable software units within AP are so called Adaptive Applications (the corresponding meta-model element is `AdaptiveAutosarApplication`). Such an Adaptive Application consists of 1..n executables, which are in turn built up by instantiating `CompositionSwComponentType` (with arbitrary nesting) as described in the previous chapter. Typically integrators then decide, which Adaptive Applications in the form of its 1..n executables they start at all and how many times they start a certain Adaptive

Application/its associated executables. That means for those kind of implicit instantiation no specific code has to be written! Integrators rather have to deal with machine configuration, to configure how many times Applications get started. A started Adaptive Application then turns into 1..n processes (depending on the number of executables it is made of). We call this then the "deployment level".

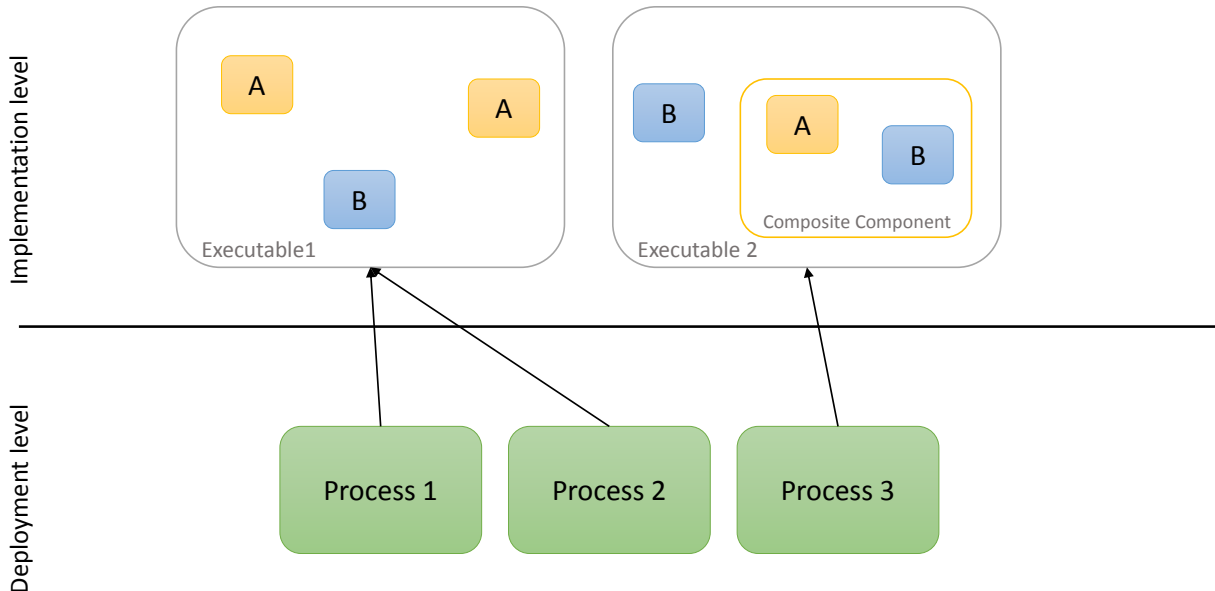


Figure 7.8: Instantiation of Adaptive Applications in Deployment

The figure above shows a simple example, where we have two Adaptive Applications, where each of those exactly consists of one executable. Adaptive Application 1 with Executable 1 is deployed twice, leading to Process 1 and Process 2 after executable start, where Application 2, which consists of Executable 2 is deployed once leading to Process 3 after start.

7.4.5 Usage of meta-model identifiers within ara::com based application code

The explanations of meta-model/ara::com relation up to this point should help to understand the structure of `instance specifiers` used in `ResolveInstanceIDs` described in Listing 4.3. As described in the previous chapter and depicted in Figure 7.6 the `instance specifiers` relate in a certain way to the corresponding port in the model of the `SoftwareComponentType`. If you followed the previous chapters the **port name** of the model alone isn't sufficient to clearly identify it in its final instantiation, where the same component implementation might be instantiated multiple times in the code and then eventually started multiple times in different processes. Instance IDs obviously have to be assigned to objects, which finally have a distinct identity in an deployment.

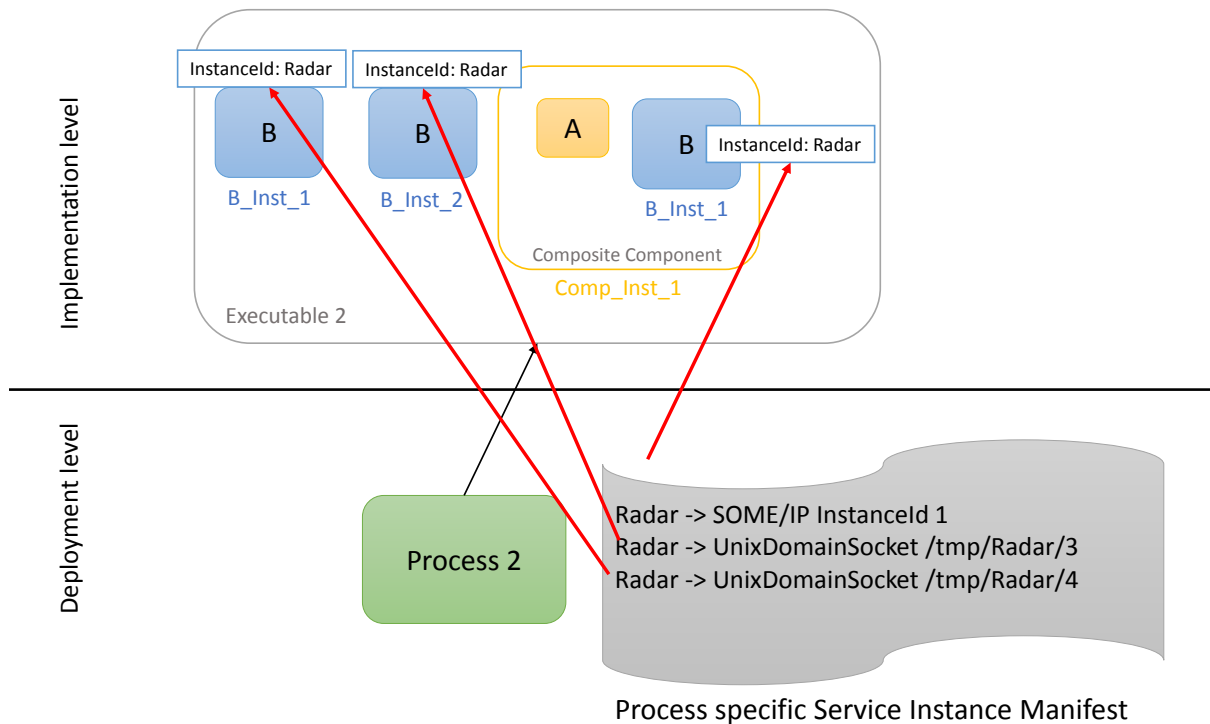


Figure 7.9: Instancelids in Deployment

The figure above outlines the "problem" with a simple example. Within Executable 2 there are three instantiations of `SoftwareComponentType B` implementation in different contexts (nesting levels). All instances do provide a specific instance of SI `RadarService`. The integrator, who applies the `Service Instance Manifest` for Process 2 has to do the technical mapping on `ara::com` level. I.e. he has to decide, which technical transport binding is to be used in each of the B instantiations and subsequently also, which technical transport binding specific instance ID. In our example, the integrator wants to provide the SI `RadarService` via `SOME/IP` binding and an `SOME/IP` specific instance ID "1" in the context of the B instantiation, which is nested inside the composite component on the right side, while he decides to provide the SI `RadarService` via local `IPC` (Unix domain socket) binding and a Unix domain socket specific instance ID `"/tmp/Radar/3"` and `"/tmp/Radar/4"` in the context of the B instantiations on the left side, which are not nested (they are instantiated at "top-level" of the executable). Here it gets obvious, that within the `Service Instance Manifest`, which allows to specify the mapping of port instantiations within a Process to technical bindings and their concrete instance IDs, the sole usage of the **port name** from the model isn't sufficient to differentiate. To get unique identifiers within an executable (and therefore a process), the nature of nested instantiation and re-use of `SoftwareComponentTypes` has to be considered. Every time a `SoftwareComponentType` gets instantiated, its instantiation gets a unique name within its instantiation context. This concept applies to both: C++ implementation level and AUTOSAR meta-model level! In our concrete example this means:

- B instantiations on top level get unique names on their level: "B_Inst_1" and "B_Inst_2"

- B instantiation within the Composite Component Type gets unique name on this level: "B_Inst_1"
- Composite Component instantiation on top level gets unique name on its level: "Comp_Inst_1"
- From the perspective of the executable/process, we therefore have unique identifiers for all instances of B:
 - "B_Inst_1"
 - "B_Inst_2"
 - "Comp_Inst_1::B_Inst_1"

For an Adaptive Software Component developer this then means in a nutshell:

If you construct an `instance specifier` to be transformed via `ResolveInstanceIDs()` into an `ara::com::InstanceIdentifier` or used directly with `FindService()` (R-port side from model perspective) or as `ctor` parameter for a skeleton (P-port side from model perspective), it shall look like:

```
<context identifier>/<port name>
```

Port name is to be taken from the model, which describes the `AdaptiveApplicationSwComponentType` to be developed. Since you are not necessarily the person who decides where and how often your component gets deployed, you should foresee, that your `AdaptiveApplicationSwComponentType` implementation can be handed over a stringified `<context identifier>`, which you

- either use directly, when constructing `ara::core::InstanceSpecifier` to instantiate proxies/skeleton, which reflect your own component ports.
- "hand over" to other `AdaptiveApplicationSwComponentType` implementations, which you instantiate from your own `AdaptiveApplicationSwComponentType` implementation (that is creating a new nesting level)

Note: Since AUTOSAR AP does **not** prescribe, how the component model on meta-model level shall be translated to (C++) implementation level, component instantiation (nesting of components) and "handing over" of the `<context identifier>` is up to the implementer! It might be a "natural" solution, to solve this by a `<context identifier>` `ctor` parameter for multi instantiable `AdaptiveApplicationSwComponentTypes`.

7.5 Abstract Protocol Network Binding Examples

This chapter presents Abstract Protocol Network Bindings examples using an `InstanceSpecifier`.

Proxy "FindService" Code Examples :




Figure 7.10: Find Service using abstract network binding

Proxy "ANY" Code Examples :

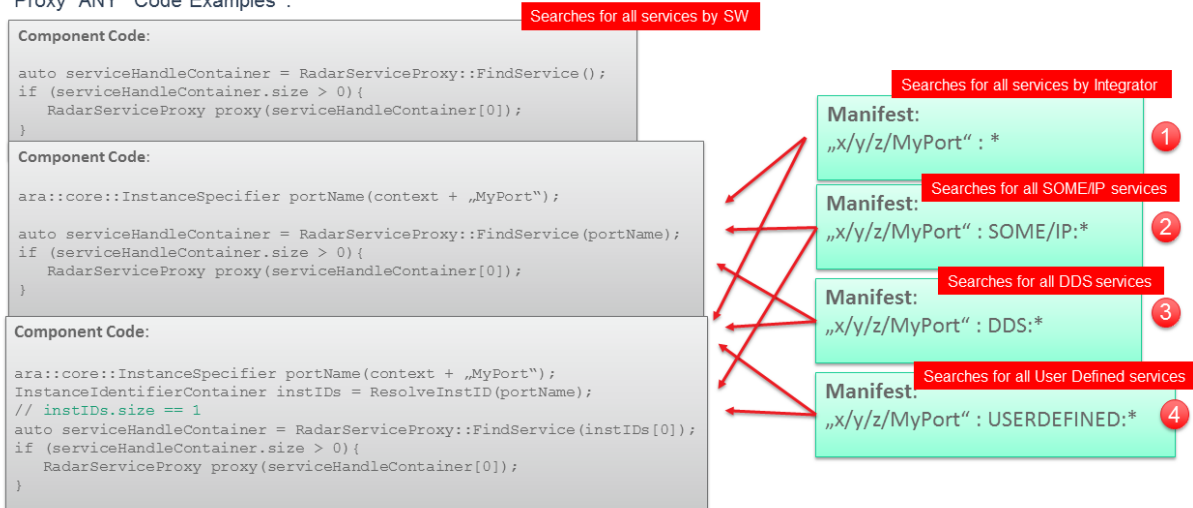


Figure 7.11: Find Service using abstract network binding - ANY

Skeleton "ctor" Code Examples :



Figure 7.12: Skeleton creation using abstract network bindings

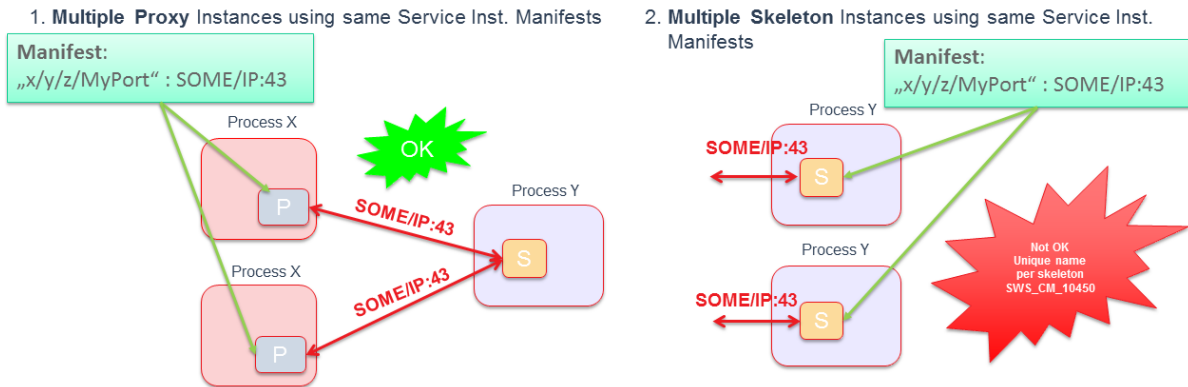


Figure 7.13: Multiple usage of the same service instance manifest for an abstract binding