

Document Title	Specification of RTE
Document Owner	AUTOSAR
Document Responsibility	AUTOSAR
Document Identification No	084
Document Classification	Standard

Document Version	2.3.0
Document Status	Final
Part of Release	3.1
Revision	5

Document Change History			
Date	Version	Changed by	Change Description
22.09.2010	2.3.0	AUTOSAR Administration	<ul style="list-style-type: none"> • Generation of the indirect API decoupled from multiple instantiation: changed rte_sws_1355, rte_sws_2613, rte_sws_2615. • Behavior in name clashes of AUTOSAR types PIM types: added rte_sws_5195, changed rte_sws_3789, rte_sws_3782.

27.01.2010	2.2.0	AUTOSAR Administration	<ul style="list-style-type: none"> • Allow Communication Attributes on Compositions: changed <code>rte_sws_in_0055</code>, <code>rte_sws_in_0062</code>, <code>rte_sws_in_5023</code>, <code>rte_sws_in_5050</code>, <code>rte_sws_in_0067</code>, <code>rte_sws_in_0029</code>, <code>rte_sws_in_2701</code>, <code>rte_sws_in_2693</code> • Support for initial calibration data values: added <code>rte_sws_7186</code>, <code>rte_sws_7185</code>, <code>rte_sws_2750</code>. • Reverted implementation of "incompatible function declarations": changed <code>rte_sws_1017</code>, <code>rte_sws_1018</code>, <code>rte_sws_1019</code>, <code>rte_sws_1020</code>, <code>rte_sws_5107</code>, <code>rte_sws_5108</code>, <code>rte_sws_5109</code>, <code>rte_sws_1254</code>, <code>rte_sws_3930</code>, <code>rte_sws_3593</code>, <code>rte_sws_5512</code>; added <code>rte_sws_5195</code>, <code>rte_sws_5196</code>, <code>rte_sws_5197</code>, <code>rte_sws_5198</code>, <code>rte_sws_5199</code>, <code>rte_sws_5200</code>, <code>rte_sws_5201</code>, <code>rte_sws_5202</code>, <code>rte_sws_5203</code>, <code>rte_sws_5204</code>, <code>rte_sws_5205</code>, <code>rte_sws_5206</code>, <code>rte_sws_5207</code>, <code>rte_sws_5208</code>, <code>rte_sws_5209</code>; removed <code>rte_sws_3743</code>; Fixed typo in <code>rte_sws_6129</code>, <code>rte_sws_3750</code> (CalPrm vs. Calprm).
04.02.2009	2.1.0	AUTOSAR Administration	<ul style="list-style-type: none"> • updated VFB-Tracing: changes <code>rte_sws_1327</code>, <code>rte_sws_1328</code> • unconnected R-Ports are supported: changed <code>rte_sws_1329</code>, <code>rte_sws_3019</code>; added <code>rte_sws_1330</code>, <code>rte_sws_1331</code>, <code>rte_sws_1333</code>, <code>rte_sws_1334</code>, <code>rte_sws_1336</code>, <code>rte_sws_1337</code>, <code>rte_sws_1346</code>, <code>rte_sws_2621</code>, <code>rte_sws_2638</code>, <code>rte_sws_2639</code>, <code>rte_sws_2640</code>, <code>rte_sws_3785</code>, <code>rte_sws_5099</code>, <code>rte_sws_5100</code>, <code>rte_sws_5101</code>, <code>rte_sws_5102</code> • incompatible function declarations: changed <code>rte_sws_1018</code>, <code>rte_sws_1019</code>, <code>rte_sws_1020</code>; added <code>rte_sws_5107</code>, <code>rte_sws_5108</code>, <code>rte_sws_5109</code>; removed <code>rte_sws_6030</code>. • Insufficient RTE server mapping requirement: changed <code>rte_sws_2204</code>.

15.02.2008	2.0.1	AUTOSAR Administration	Layout adaptations
20.12.2007	2.0.0	AUTOSAR Administration	<ul style="list-style-type: none"> • Adapted to new version of meta model • "RTE ECU Configuration" added • Calibration and measurement revised • Document meta information extended • Small layout adaptations made
31.01.2007	1.1.1	AUTOSAR Administration	<ul style="list-style-type: none"> • "Advice for users" revised • "Revision Information" added
01.12.2006	1.1.0	AUTOSAR Administration	<p>Updated for AUTOSAR Release 2.1.</p> <ul style="list-style-type: none"> • Adapted to new version of meta model • New feature 'debouncing of runnable activation' • New feature 'runnable activation offset' • 'Measurement and Calibration' added • Semantics of implicit communication enhanced • Legal disclaimer revised
18.07.2006	1.0.1	AUTOSAR Administration	Second release. Additional features integrated, adapted to updated version of meta-model.
05.05.2006	1.0.0	AUTOSAR Administration	Initial release

Disclaimer

This specification 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 specification.

The material contained in this specification is protected by copyright and other types of Intellectual Property Rights. The commercial exploitation of the material contained in this specification requires a license to such Intellectual Property Rights.

This specification 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 specification may be utilized or reproduced, in any form or by any means, without permission in writing from the publisher.

The AUTOSAR specifications have been developed for automotive applications only. They have neither been developed, nor tested for non-automotive applications.

The word AUTOSAR and the AUTOSAR logo are registered trademarks.

Advice for users

AUTOSAR Specification Documents may contain exemplary items (exemplary reference models, "use cases", and/or references to exemplary technical solutions, devices, processes or software).

Any such exemplary items are contained in the Specification Documents for illustration purposes only, and they themselves are not part of the AUTOSAR Standard. Neither their presence in such Specification Documents, nor any later documentation of AUTOSAR conformance of products actually implementing such exemplary items, imply that intellectual property rights covering such exemplary items are licensed under the same rules as applicable to the AUTOSAR Standard.

Table of Contents

1	Introduction	14
1.1	Scope	14
1.2	Dependency to other AUTOSAR specifications	15
1.3	Acronyms and Abbreviations	16
1.4	Technical Terms	16
1.5	Document Conventions	17
1.6	Requirements Traceability	18
2	RTE Overview	28
2.1	The RTE in the Context of AUTOSAR	28
2.2	AUTOSAR Concepts	28
2.2.1	AUTOSAR Software-components	28
2.2.2	Basic Software Modules	29
2.2.3	Communication	29
2.2.3.1	Communication Models	29
2.2.3.2	Communication Modes	30
2.2.3.3	Static Communication	30
2.2.3.4	Multiplicity	30
2.2.4	Concurrency	31
2.3	The RTE Generator	32
2.4	Design Decisions	32
3	RTE Generation Process	34
3.1	RTE Contract Phase	36
3.2	RTE Configuration Editing	37
3.3	RTE Generation Phase	38
3.4	RTE Configuration and OS Interaction	38
4	RTE Functional Specification	41
4.1	Architectural concepts	41
4.1.1	Scope	41
4.1.2	RTE and AUTOSAR Software-Components	41
4.1.2.1	Structure of SW-Components	42
4.1.2.2	Ports, Interfaces and Connections	42
4.1.2.3	Internal Behavior	44
4.1.2.4	Implementation	46
4.1.3	Instantiation	47
4.1.3.1	Scope and background	47
4.1.3.2	Concepts of instantiation	48
4.1.3.3	Single instantiation	49
4.1.3.4	Multiple instantiation	49
4.1.4	RTE and AUTOSAR Services	50
4.1.5	RTE and ECU Abstraction	51
4.1.6	RTE and Complex Device Driver	52

4.2	RTE Implementation Aspects	52
4.2.1	Scope	52
4.2.2	OS	54
4.2.2.1	OS Objects	54
4.2.2.2	Runnable Entities	56
4.2.2.3	RTE Events	57
4.2.2.4	Mapping of runnable entities to tasks	58
4.2.2.5	Activation Offset for runnable	66
4.2.2.6	Activation and Start of Runnable Entities	68
4.2.3	Interrupt decoupling and notifications	71
4.2.3.1	Basic notification principles	71
4.2.3.2	Interrupts	72
4.2.3.3	Decoupling interrupts on RTE level	73
4.2.3.4	RTE and interrupt categories	74
4.2.4	Data Consistency	74
4.2.4.1	General	74
4.2.4.2	Communications to look at	76
4.2.4.3	Concepts	76
4.2.4.4	Mechanisms to guarantee data consistency	77
4.2.4.5	Exclusive Areas	80
4.2.4.6	InterRunnableVariables	82
4.2.5	Multiple trigger of Runnables	85
4.2.6	Measurement and Calibration	86
4.2.6.1	General	86
4.2.6.2	Measurement	87
4.2.6.3	Calibration	92
4.3	Communication Models	107
4.3.1	Sender-Receiver	107
4.3.1.1	Introduction	107
4.3.1.2	Receive Modes	108
4.3.1.3	Multiple Data Elements	110
4.3.1.4	Multiple Receivers and Senders	112
4.3.1.5	Implicit and Explicit Data Reception and Transmission	112
4.3.1.6	Transmission Acknowledgement	117
4.3.1.7	Communication Time-out	119
4.3.1.8	Data Element Invalidation	120
4.3.1.9	Filters	122
4.3.1.10	Buffering	122
4.3.1.11	Operation	123
4.3.2	Client-Server	131
4.3.2.1	Introduction	131
4.3.2.2	Multiplicity	132
4.3.2.3	Communication Time-out	134
4.3.2.4	Port-Defined argument values	135
4.3.2.5	Buffering	136
4.3.2.6	Inter ECU Response to Request Mapping	137

4.3.2.7	Operation	139
4.3.3	SWC internal communication	144
4.3.3.1	InterRunnableVariables	144
4.4	Modes	145
4.4.1	Mode User	145
4.4.2	Mode Manager	147
4.4.3	Refinement of the semantics of ModeDeclarations and Mode- DeclarationGroups	148
4.4.4	Order of actions taken by the RTE upon interception of a mode switch notification	148
4.4.5	Notification of mode switches	152
4.5	Initialization and Finalization	155
4.5.1	Initialization and Finalization of the RTE	155
4.5.2	Initialization and Finalization of AUTOSAR Software-Components	155
4.6	RTE Functionality Levels	156
5	RTE Reference	157
5.1	Scope	157
5.1.1	Programming Languages	157
5.1.2	Generator Principles	158
5.1.2.1	Operating Modes	158
5.1.2.2	Optimization Modes	159
5.1.3	Generator external configuration switches	160
5.2	API Principles	160
5.2.1	RTE Namespace	161
5.2.2	Direct API	161
5.2.3	Indirect API	162
5.2.3.1	Accessing Port Handles	162
5.2.4	DataReadAccess and DataWriteAccess	163
5.2.5	PerInstanceMemory	164
5.2.6	API Mapping	166
5.2.6.1	“RTE Contract” Phase	167
5.2.6.2	“RTE Generation” Phase	168
5.2.6.3	Function Elidation	168
5.2.6.4	API Naming Conventions	169
5.2.6.5	API Parameters	170
5.2.6.6	Error Handling	174
5.2.6.7	Success Feedback	175
5.2.7	Unconnected Ports	175
5.2.7.1	Data Elements	176
5.2.7.2	Mode Ports	177
5.2.7.3	Client-Server	178
5.2.8	Non-identical ports	178
5.3	RTE Modules	179
5.3.1	RTE Header File	179
5.3.2	Lifecycle Header File	180

5.3.3	Application Header File	180
5.3.3.1	File Name	181
5.3.3.2	Scope	181
5.3.3.3	File Contents	182
5.3.4	AUTOSAR Types Header File	184
5.3.4.1	File Contents	184
5.3.4.2	Primitive AUTOSAR Data Types	184
5.3.4.3	Complex AUTOSAR Data Types	186
5.3.4.4	C/C++	188
5.3.5	VFB Tracing Header File	188
5.3.5.1	C/C++	188
5.3.5.2	File Contents	188
5.3.6	RTE Configuration Header File	190
5.3.6.1	C/C++	190
5.3.6.2	File Contents	190
5.3.7	Generated RTE	191
5.3.7.1	Header File Usage	191
5.3.7.2	C/C++	192
5.3.7.3	File Contents	192
5.3.7.4	Reentrancy	194
5.4	RTE Data Structures	194
5.4.1	Instance Handle	194
5.4.2	Component Data Structure	195
5.4.2.1	Data Handles Section	197
5.4.2.2	Per-instance Memory Handles Section	199
5.4.2.3	Inter Runnable Variable Handles Section	200
5.4.2.4	Exclusive-area handles Section	200
5.4.2.5	Port API Section	201
5.4.2.6	Calibration Parameter Handles Section	205
5.4.2.7	Inter Runnable Variable API Section	205
5.4.2.8	Vendor Specific Section	206
5.5	API Data Types	206
5.5.1	Std_ReturnType	206
5.5.1.1	Infrastructure Errors	208
5.5.1.2	Application Errors	208
5.5.1.3	Predefined Error Codes	209
5.5.2	Rte_Instance	211
5.5.3	RTE Modes	211
5.5.4	Enumeration Data Types	212
5.5.5	Range Data Types	213
5.6	API Reference	213
5.6.1	Rte_Ports	214
5.6.2	Rte_NPorts	214
5.6.3	Rte_Port	215
5.6.4	Rte_Send/Rte_Write	215
5.6.5	Rte_Switch	217

5.6.6	Rte_Invalidate	218
5.6.7	Rte_Feedback	219
5.6.8	Rte_Read	222
5.6.9	Rte_Receive	223
5.6.10	Rte_Call	224
5.6.11	Rte_Result	226
5.6.12	Rte_Pim	228
5.6.13	Rte_CData	229
5.6.14	Rte_Calprm	229
5.6.15	Rte_IRead	230
5.6.16	Rte_IWrite	231
5.6.17	Rte_IWriteRef	232
5.6.18	Rte_IInvalidate	233
5.6.19	Rte_IStatus	234
5.6.20	Rte_IrvIRead	235
5.6.21	Rte_IrvIWrite	236
5.6.22	Rte_IrvRead	237
5.6.23	Rte_IrvWrite	238
5.6.24	Rte_Enter	238
5.6.25	Rte_Exit	239
5.6.26	Rte_Mode	240
5.7	Runnable Entity Reference	241
5.7.1	Signature	241
5.7.2	Entry Point Prototype	241
5.7.3	Role Parameters	242
5.7.4	Return Value	242
5.7.5	Triggering Events	242
5.7.5.1	TimingEvent	243
5.7.5.2	ModeSwitchEvent	243
5.7.5.3	AsynchronousServerCallReturnsEvent	243
5.7.5.4	DataReceiveErrorEvent	243
5.7.5.5	OperationInvokedEvent	244
5.7.5.6	DataReceivedEvent	245
5.7.5.7	DataSendCompletedEvent	245
5.7.6	Reentrancy	245
5.8	RTE Lifecycle API Reference	246
5.8.1	Rte_Start	246
5.8.2	Rte_Stop	247
5.9	RTE Call-backs Reference	248
5.9.1	RTE-COM Message Naming Conventions	248
5.9.2	Communication Service Call-backs	248
5.9.3	Naming convention of CallbackRoutineName	249
5.10	VFB Tracing Reference	250
5.10.1	Prinicple of Operation	251
5.10.2	Trace Events	251
5.10.2.1	RTE API Trace Events	251

5.10.2.2	COM Trace Events	252
5.10.2.3	OS Trace Events	253
5.10.2.4	Runnable Entity Trace Events	255
5.10.3	Configuration	256
5.10.4	Interaction with Object-code Software-Components	256
6	RTE ECU Configuration	257
6.1	RTE Generation Parameters	260
6.2	Handling of Software Component instances	263
6.2.1	Selection of SW-Component Implementation	264
6.2.2	Runnable Entity to task mapping	265
6.2.3	Exclusive Area implementation	268
6.2.4	NVRam Allocation	270
6.3	Component Type Calibration	272
6.4	Communication infrastructure	273
A	Metamodel Restrictions	274
A.1	Restriction concerning WaitPoint	274
A.2	Restriction concerning RTEEvent	274
A.3	Restriction concerning isQueued attribute of DataElementPrototype	275
A.4	Restriction concerning ServerCallPoint	275
A.5	Restriction concerning multiple instantiation of software components	276
A.6	Restriction concerning runnable entity	276
A.7	Restrictions concerning runnables with dependencies on modes	276
A.8	Restriction concerning InterRunnableVariables	278
A.9	Restriction concerning InternalBehavior	278
A.10	Restriction concerning Initial Value	279
A.11	Restriction concerning PerInstanceMemory	279
A.12	Restriction concerning unconnected r-port	279
A.13	Restrictions regarding n:1 sender-receiver communication	280
A.14	Restrictions regarding Measurement and Calibration	280
A.15	Restriction concerning ExclusiveAreaImplMechanism	280
B	Required Input Information	282
B.1	SWC and instance	282
B.2	Runnable entity and task	286
B.3	Port and interface	291
B.4	Communication	299
B.5	Data consistency	301
B.6	RTE configuration	303
B.7	Measurement and calibration	305
B.8	Mode management	306
C	External Requirements	310
D	MISRA C Compliance	312

E Interfaces of COM used by the RTE

313

Bibliography

- [1] Glossary
AUTOSAR_Glossary.pdf
- [2] Methodology
AUTOSAR_Methodology.pdf
- [3] Requirements on Communication
AUTOSAR_SRS_COM.pdf
- [4] Requirements on ECU Configuration
AUTOSAR_RS_ECU_Configuration.pdf
- [5] Requirements on Operating System
AUTOSAR_SRS_OS.pdf
- [6] Specification of Communication
AUTOSAR_SWS_COM.pdf
- [7] Specification of ECU Configuration
AUTOSAR_ECU_Configuration.pdf
- [8] Specification of ECU State Manager
AUTOSAR_SWS_ECU_StateManager.pdf
- [9] Specification of Interoperability of Authoring Tools
AUTOSAR_InteroperabilityAuthoringTools.pdf
- [10] Specification of I/O Hardware Abstraction
AUTOSAR_SWS_IO_HWAbstraction.pdf
- [11] Specification of Operating System
AUTOSAR_SWS_OS.pdf
- [12] Specification of Standard Types
AUTOSAR_SWS_StandardTypes.pdf
- [13] Specification of the Virtual Functional Bus

AUTOSAR_VirtualFunctionBus.pdf

[14] Specification of System Template
AUTOSAR_SystemTemplate.pdf

[15] DTD File
AUTOSAR_DTD_File.dtd

[16] Template Modeling Guide
AUTOSAR_TemplateModelingGuide.pdf

[17] Software Component Template
AUTOSAR_SoftwareComponentTemplate.pdf

[18] Gemeinsames Subset der MISRA C Guidelines
[http://www.automotive-his.de/download/
HIS_SubSet_MISRA_C_1.0.3.pdf](http://www.automotive-his.de/download/HIS_SubSet_MISRA_C_1.0.3.pdf)

Note on XML examples

This specification includes examples in XML based on the AUTOSAR metamodel available at the time of writing. These examples are included as illustrations of configurations and their expected outcome but should not be considered part of the specification.

1 Introduction

This document contains the software specification of the AUTOSAR Run-Time Environment (*RTE*). Basically, the RTE together with the OS, AUTOSAR COM and other Basic Software Modules is the implementation of the Virtual Functional Bus concepts (*VFB*, [13]). The RTE implements the AUTOSAR Virtual Functional Bus interfaces and thereby realizes the communication between AUTOSAR software-components.

This document describes how these concepts are realized within the RTE. Furthermore, the Application Programming Interface (*API*) of the RTE and the interaction of the RTE with other basic software modules is specified.

1.1 Scope

This document is intended to be the main reference for developers of an RTE generator tool or of a concrete RTE implementation respectively. The document is also the reference for developers of AUTOSAR software-components and basic software modules that interact with the RTE, since it specifies the application programming interface of the RTE and therefore the mechanisms for accessing the RTE functionality. Furthermore, this specification should be read by the AUTOSAR working groups that are closely related to the RTE (see Section 1.2 below), since it describes the interfaces of the RTE to these modules as well as the behavior / functionality the RTE expects from them.

This document is structured as follows. After this general introduction, Chapter 2 gives a more detailed introduction of the concepts of the RTE. Chapter 3 describes how an RTE is generated in the context of the overall AUTOSAR methodology. Chapter 4 is the central part of this document. It specifies the RTE functionality in detail. The RTE API is described in Chapter 5.

The appendix of this document consists of five parts: Appendix A lists the restrictions to the AUTOSAR metamodel that this version of the RTE specification relies on. Appendix B describes the input that is needed for the RTE generation process and where this input is assumed to come from. Appendix C explicitly lists all external requirements, i.e. all requirements that are not about the RTE itself but specify the assumptions on the environment and the input of an RTE generator. In Appendix D some HIS MISRA rules are listed that are likely to be violated by RTE code, and the rationale why these violations may occur. Finally, Appendix E lists the COM API and COM Callback functions that are used by the RTE

Note that Chapters 1 and 2, as well as Appendix D and E do not contain any requirements and are thus intended for information only.

Chapters 4, 5, and Appendix B are probably of most interest for developers of an RTE Generator. Chapters 2, 3, 5 are important for developers of AUTOSAR software-components and basic software modules. The most important chapters for related AUTOSAR work packages would be Chapters 4 and 5, as well as Appendix B and C.

The specifications in this document do not define details of the implementation of a concrete RTE or RTE generator respectively. Furthermore, aspects of the ECU- and system-generation process (like e.g. the mapping of SW-Cs to ECUs, or schedulability analysis) are also not in the scope of this specification. Nevertheless, it is specified what input the RTE generator expects from these configuration phases.

1.2 Dependency to other AUTOSAR specifications

The main documents that served as input for the specification of the RTE are the specification of the Virtual Functional Bus [13] and the specification of the Software Component Template [17]. Also of primary importance are the specifications of those Basic Software modules that closely interact with the RTE (or vice versa). These are especially the communication module [6] and the operating system [11]. The main input of an RTE generator is described (among others) in the ECU Configuration Description. Therefore, the corresponding specification [4] is also important for the RTE specification. Furthermore, as the process of RTE generation is an important part of the overall AUTOSAR Methodology, the corresponding document [2] is also considered.

The following list shows the specifications that are closely interdependent to the specification of the RTE:

- Specification of the Virtual Functional Bus [13]
- Specification of the Software Component Template [17]
- Specification of AUTOSAR COM [6]
- Specification of AUTOSAR OS [11]
- Specification of ECU State Manager and Communication Manager [8]
- Specification of ECU-Configuration Description / Generation [4]
- Specification of System Description / Generation [14]
- AUTOSAR Methodology [2]
- Documents relevant for the AUTOSAR Metamodel [16, 15]

1.3 Acronyms and Abbreviations

All abbreviations used throughout this document – except the ones listed here – can be found in the official AUTOSAR glossary [1].

1.4 Technical Terms

All technical terms used throughout this document – except the ones listed here – can be found in the official AUTOSAR glossary [1] or the Software Component Template Specification [17].

Term	Description
mode switch interface	A <code>SenderReceiverInterface</code> with a <code>ModeDeclarationGroupPrototype</code> of a <code>ModeDeclarationGroup</code> is called mode switch interface for the <code>ModeDeclarationGroup</code> . The <code>mode ports</code> of the mode manager and the <code>mode user</code> are of the type of a mode switch interface. Beware, a <code>SenderReceiverInterface</code> may contain any combination of <code>DataElementPrototypes</code> and <code>ModeDeclarationGroupPrototypes</code> .
mode port	The port for receiving (or sending) a mode switch notification. For this purpose, a mode port is typed by a <code>mode switch interface</code> .
mode user	An AUTOSAR SW-C that depends on modes by <code>ModeDisablingDependency</code> , <code>ModeSwitchEvent</code> , or simply by reading the current state of a mode is called a <code>mode user</code> . A mode user is defined by having a <code>require mode port</code> . See also section 4.4.1.
mode manager	Entering and leaving modes is initiated by a mode manager. A mode manager is defined by having a <code>provide mode port</code> . A mode manager might be either an <code>application mode manager</code> or a basic software module that provides a service including mode switches, like the ECU State Manager. See also section 4.4.2.
application mode manager	An <code>application mode manager</code> is an AUTOSAR Software-Component that provides the service of switching modes. The modes of an <code>application mode manager</code> do not have to be standardized.
mode switch notification	The communication of a mode switch from the mode manager to the mode user using the <code>mode switch interface</code> is called mode switch notification.

mode machine instance	<p>The instances of mode machines or ModeDeclarationGroups are defined by the ModeDeclarationGroupPrototypes of the <code>mode managers</code>.</p> <p>Since a mode switch is not executed instantaneously, The RTE has to maintain it's own states. For each <code>mode manager</code>'s ModeDeclarationGroupPrototype, RTE has one state machine. This state machine is called mode machine instance. For all <code>mode users</code> of the same <code>mode manager</code>'s ModeDeclarationGroupPrototype, RTE uses the same mode machine instance. See also section 4.4.2.</p>
mode disabling dependent runnable	<p>A mode disabling dependent runnable is triggered by an RteEvent with a ModeDisablingDependency. RTE prevents the start of that runnable by the RteEvent, when the corresponding <code>mode disabling</code> is active. See also section 4.4.1.</p>
mode disabling	<p>When a 'mode disabling' is active, RTE disables the start of <code>mode disabling dependent runnables</code>. The 'mode disabling' is active during the mode that is referenced in the mode disabling dependency and during the transitions that enter and leave this mode. See also section 4.4.1.</p>
OnEntry runnable	<p>A runnable that is triggered by a ModeSwitchEvent with ModeActivationKind 'entry' is triggered on entering the mode. It is called OnEntry runnable. See also section 4.4.1.</p>
OnExit runnable	<p>A runnable that is triggered by a ModeSwitchEvent with ModeActivationKind 'exit' is triggered on exiting the mode. It is called OnExit runnable. See also section 4.4.1.</p>
server runnable	<p>A server that is triggered by an OperationInvokedEvent. It has a mixed behaviour between a runnable and a function call. In certain situations, RTE can implement the client server communication as a simple function call.</p>
runnable activation	<p>The activation of a runnable is defined as the activation of the task that contains the runnable and eventually includes setting a flag that tells the glue code in the task which runnable is to be executed.</p>
runnable start	<p>A runnable is started by the calling the C-function that implements the runnable from within a started task.</p>

1.5 Document Conventions

Requirements in the SRS are referenced using [RTE<n>] where <n> is the requirement id. For example, [RTE00098].

Requirements in the SWS are marked with **[rte_sws_<n>]** as the first text in a paragraph. The scope of the requirement is the entire paragraph.

Requirements on the input of the RTE specified in terms of the meta model are marked with **[rte_sws_in_<n>]** accordingly.

External requirements on the input of the RTE are marked with **[rte_sws_ext_<n>]**.

Technical terms are typeset in monospace font, e.g. `Warp Core`.

API function calls are also marked with monospace font, like `Rte_ejectWarpCore()`.

1.6 Requirements Traceability

Requirement	Satisfied by
[BSW00300] Module naming convention	rte_sws_1171 rte_sws_1157 rte_sws_1158 rte_sws_1003 rte_sws_1161 rte_sws_1169
[BSW00304] AUTOSAR integer data types	rte_sws_1175 rte_sws_1176 rte_sws_1177 rte_sws_1178 rte_sws_1179 rte_sws_1180 rte_sws_1181 rte_sws_1182 rte_sws_1183 rte_sws_1184 rte_sws_1185
[BSW00305] Self-defined data types naming convention	rte_sws_1150 rte_sws_3713 rte_sws_3714 rte_sws_3733 rte_sws_2301 rte_sws_3731 rte_sws_1055
[BSW00307] Global variables naming convention	rte_sws_1171 rte_sws_3712
[BSW00308] Definition of global data	not testable
[BSW00310] API naming convention	rte_sws_1071 rte_sws_1072 rte_sws_2631 rte_sws_1206 rte_sws_1083 rte_sws_1091 rte_sws_1092 rte_sws_1102 rte_sws_1111 rte_sws_1118 rte_sws_1252 rte_sws_3928 rte_sws_3741 rte_sws_3744 rte_sws_5509 rte_sws_3800 rte_sws_3550 rte_sws_3553 rte_sws_3560 rte_sws_3565 rte_sws_1120 rte_sws_1123 rte_sws_2569
[BSW00312] Shared code shall be reentrant	rte_sws_3749
[BSW00326] Transition from ISRs to OS tasks	rte_sws_3600 rte_sws_3594 rte_sws_3530 rte_sws_3531 rte_sws_3532
[BSW00327] Error values naming convention	rte_sws_1058 rte_sws_1060 rte_sws_1064 rte_sws_1317 rte_sws_1061 rte_sws_1065 rte_sws_2571
[BSW00330] Usage of macros / inline functions instead of functions	rte_sws_1274
[BSW007] HIS MISRA C	rte_sws_3715 rte_sws_1168
[RTE00003] Tracing of sender-receiver communication	rte_sws_1357 rte_sws_1238 rte_sws_1240 rte_sws_1241 rte_sws_3814 rte_sws_1242

[RTE00004] Tracing of client-server communication	rte_sws_1357 rte_sws_1238 rte_sws_1240 rte_sws_1241 rte_sws_3814 rte_sws_1242
[RTE00005] Support for 'trace' build	rte_sws_3607 rte_sws_1320 rte_sws_1322 rte_sws_1323 rte_sws_1327 rte_sws_1328
[RTE00008] VFB tracing configuration	rte_sws_3607 rte_sws_1320 rte_sws_1236 rte_sws_1321 rte_sws_1322 rte_sws_1323 rte_sws_1324 rte_sws_1325
[RTE00011] Support for multiple AUTOSAR software-component instances	rte_sws_2000 rte_sws_2001 rte_sws_2008 rte_sws_2009 rte_sws_2002 rte_sws_2017 rte_sws_1148 rte_sws_1012 rte_sws_1013 rte_sws_3806 rte_sws_3793 rte_sws_3713 rte_sws_3718 rte_sws_3719 rte_sws_1349 rte_sws_3720 rte_sws_3721 rte_sws_3716 rte_sws_3717 rte_sws_3722 rte_sws_3711 rte_sws_1016
[RTE00012] Multiply instantiated AUTOSAR software-components delivered as binary code shall share code	rte_sws_3015 rte_sws_2017 rte_sws_1007
[RTE00013] Static memory sections	rte_sws_3790 rte_sws_2303 rte_sws_2304 rte_sws_3789 rte_sws_3782 rte_sws_2305 rte_sws_5062 rte_sws_2301 rte_sws_2302
[RTE00017] Rejection of inconsistent component implementations	rte_sws_3755 rte_sws_4504 rte_sws_3764 rte_sws_1004 rte_sws_1276
[RTE00018] Rejection of invalid configurations	rte_sws_2750 rte_sws_5508 rte_sws_7006 rte_sws_5508 rte_sws_2254 rte_sws_2100 rte_sws_2051 rte_sws_2009 rte_sws_2204 rte_sws_1313
[RTE00019] RTE is the communication infrastructure	rte_sws_6000 rte_sws_6011 rte_sws_5500 rte_sws_4527 rte_sws_6023 rte_sws_4526 rte_sws_6024 rte_sws_3760 rte_sws_3761 rte_sws_3762 rte_sws_4515 rte_sws_4516 rte_sws_4520 rte_sws_4522 rte_sws_2527 rte_sws_2528 rte_sws_3769 rte_sws_1048 rte_sws_1231 rte_sws_5063 rte_sws_3007 rte_sws_3008 rte_sws_3000 rte_sws_3001 rte_sws_3002 rte_sws_3775 rte_sws_2612 rte_sws_2610 rte_sws_3004 rte_sws_3005 rte_sws_3776 rte_sws_5065 rte_sws_2611 rte_sws_1264 rte_sws_3795 rte_sws_3796
[RTE00020] Access to OS	rte_sws_2250 rte_sws_5071 rte_sws_5076 rte_sws_5077 rte_sws_5072
[RTE00021] Per-ECU RTE customization	rte_sws_5000 rte_sws_1316
[RTE00022] Interaction with call-backs	rte_sws_1165
[RTE00023] RTE Overheads	rte_sws_5053
[RTE00024] Source-code AUTOSAR software components	rte_sws_1315 rte_sws_1000 rte_sws_1195
[RTE00025] Static communication	rte_sws_6026

[RTE00027] VFB to RTE mapping shall be semantic preserving	rte_sws_2200 rte_sws_2201 rte_sws_1274
[RTE00028] 1:n Sender-receiver communication	rte_sws_6023 rte_sws_4526 rte_sws_6024 rte_sws_1071 rte_sws_1072 rte_sws_1077 rte_sws_1081 rte_sws_2633 rte_sws_2635 rte_sws_1082 rte_sws_2631 rte_sws_1091 rte_sws_1092 rte_sws_1135
[RTE00029] n:1 Client-server communication	rte_sws_6019 rte_sws_4519 rte_sws_4517 rte_sws_3763 rte_sws_3770 rte_sws_3767 rte_sws_3768 rte_sws_2579 rte_sws_3769 rte_sws_1102 rte_sws_1109 rte_sws_1133 rte_sws_1359 rte_sws_1166
[RTE00031] Multiple runnable entities	rte_sws_2202 rte_sws_1126 rte_sws_1132 rte_sws_1016 rte_sws_1130 rte_sws_3749
[RTE00032] Data consistency mechanisms	rte_sws_3811 rte_sws_3514 rte_sws_3500 rte_sws_3504 rte_sws_3595 rte_sws_3596 rte_sws_3503 rte_sws_7005 rte_sws_3516 rte_sws_3517 rte_sws_3519 rte_sws_1122 rte_sws_3739 rte_sws_3740 rte_sws_3812
[RTE00033] Serialization of server runnables	rte_sws_4515 rte_sws_4518 rte_sws_4522 rte_sws_2527 rte_sws_2528 rte_sws_2529 rte_sws_2530 rte_sws_7008 rte_sws_2699
[RTE00036] Assignment to OS Applications	protection is cancelled for release 3.0
[RTE00037] The RTE shall be able to invoke functions across protection boundaries	protection is cancelled for release 3.0
[RTE00045] Standardized VFB tracing interface	rte_sws_1319 rte_sws_1250 rte_sws_1251 rte_sws_1321 rte_sws_1326 rte_sws_1238 rte_sws_1239 rte_sws_1240 rte_sws_1241 rte_sws_3814 rte_sws_1242 rte_sws_1243 rte_sws_1244 rte_sws_1245 rte_sws_1246 rte_sws_1247 rte_sws_1248 rte_sws_1249
[RTE00046] Support for 'runnable runs inside' exclusive areas	rte_sws_3500 rte_sws_3515 rte_sws_1120 rte_sws_1122 rte_sws_1123
[RTE00048] RTE Generator input	rte_sws_5001 rte_sws_5076 rte_sws_5077
[RTE00049] Construction of task bodies	rte_sws_2251 rte_sws_2254 rte_sws_2204

<p>[RTE00051] RTE API mapping</p>	<p>rte_sws_3014 rte_sws_3921 rte_sws_1269 rte_sws_1148 rte_sws_3706 rte_sws_3707 rte_sws_1143 rte_sws_1348 rte_sws_1155 rte_sws_1156 rte_sws_1153 rte_sws_1146 rte_sws_2619 rte_sws_2613 rte_sws_3602 rte_sws_2614 rte_sws_2615 rte_sws_3603 rte_sws_1354 rte_sws_1355 rte_sws_1280 rte_sws_1281 rte_sws_2632 rte_sws_1282 rte_sws_1283 rte_sws_2676 rte_sws_1284 rte_sws_2677 rte_sws_1285 rte_sws_2678 rte_sws_1286 rte_sws_2679 rte_sws_1287 rte_sws_1289 rte_sws_1291 rte_sws_1292 rte_sws_1313 rte_sws_1288 rte_sws_1290 rte_sws_1293 rte_sws_1294 rte_sws_1296 rte_sws_1297 rte_sws_1298 rte_sws_1312 rte_sws_1299 rte_sws_1119 rte_sws_1300 rte_sws_1254 rte_sws_3927 rte_sws_3952 rte_sws_3929 rte_sws_3930 rte_sws_1301 rte_sws_1268 rte_sws_3593 rte_sws_1302 rte_sws_3746 rte_sws_3747 rte_sws_5510 rte_sws_5511 rte_sws_5512 rte_sws_3801 rte_sws_1303 rte_sws_3552 rte_sws_1304 rte_sws_3557 rte_sws_3559 rte_sws_3555 rte_sws_1305 rte_sws_3562 rte_sws_3563 rte_sws_3564 rte_sws_1306 rte_sws_3567 rte_sws_3568 rte_sws_1307 rte_sws_1123 rte_sws_1308 rte_sws_3718 rte_sws_3719 rte_sws_1349 rte_sws_3720 rte_sws_3721 rte_sws_3716 rte_sws_3717 rte_sws_3723 rte_sws_3733 rte_sws_2608 rte_sws_2588 rte_sws_1363 rte_sws_1364 rte_sws_2607 rte_sws_1365 rte_sws_1366 rte_sws_3734 rte_sws_2666 rte_sws_2589 rte_sws_1367 rte_sws_2301 rte_sws_2302 rte_sws_3739 rte_sws_3740 rte_sws_3812 rte_sws_2616 rte_sws_2617 rte_sws_3799 rte_sws_3731 rte_sws_3732 rte_sws_3601 rte_sws_3730 rte_sws_2620 rte_sws_2621 rte_sws_1055 rte_sws_3726 rte_sws_2618 rte_sws_1343 rte_sws_1342 rte_sws_1053 rte_sws_6029 rte_sws_3949 rte_sws_3725 rte_sws_3752 rte_sws_2623 rte_sws_3791 rte_sws_1309 rte_sws_1310 rte_sws_1159 rte_sws_1009 rte_sws_1276 rte_sws_1266 rte_sws_1197 rte_sws_1132</p>
<p>[RTE00052] Initialization and finalization of components</p>	<p>rte_sws_2503 rte_sws_2562 rte_sws_2564</p>

[RTE00053] AUTOSAR data types	rte_sws_1282 rte_sws_3559 rte_sws_3564 rte_sws_1160 rte_sws_2648 rte_sws_1163 rte_sws_1175 rte_sws_1176 rte_sws_1177 rte_sws_1178 rte_sws_1179 rte_sws_1180 rte_sws_1181 rte_sws_1182 rte_sws_1183 rte_sws_1184 rte_sws_1185 rte_sws_1186 rte_sws_1187 rte_sws_1188 rte_sws_1265 rte_sws_1214 rte_sws_1189 rte_sws_1190 rte_sws_1191 rte_sws_1192 rte_sws_1161 rte_sws_1162
[RTE00055] Use of global namespace	rte_sws_1171
[RTE00056] Pre-defined primitive data types cannot be redefined	rte_sws_1263
[RTE00059] RTE API passes 'in' primitive data types by value	rte_sws_1017
[RTE00060] RTE API shall pass 'in' complex data types by reference	rte_sws_1018 rte_sws_5107 rte_sws_5201 rte_sws_5202
[RTE00061] 'in/out' and 'out' parameters	rte_sws_1019 rte_sws_1020 rte_sws_5196 rte_sws_5197 rte_sws_5108 rte_sws_5203 rte_sws_5204 rte_sws_5109 rte_sws_5205 rte_sws_5206
[RTE00062] Local access to basic software components	rte_sws_2100 rte_sws_2051
[RTE00064] AUTOSAR Methodology	rte_sws_5071 rte_sws_5076 rte_sws_5077 rte_sws_5072
[RTE00065] Deterministic generation	rte_sws_2514 rte_sws_5071 rte_sws_5076 rte_sws_5077 rte_sws_5072
[RTE00068] Signal initial values	rte_sws_2517 rte_sws_5078
[RTE00069] Communication timeouts	rte_sws_6002 rte_sws_6013 rte_sws_3754 rte_sws_3758 rte_sws_3759 rte_sws_3763 rte_sws_3770 rte_sws_3773 rte_sws_3771 rte_sws_3772 rte_sws_3767 rte_sws_3768 rte_sws_1064 rte_sws_1095 rte_sws_1107 rte_sws_1114
[RTE00070] Invocation order of runnables	rte_sws_2207
[RTE00072] Activation of runnable entities	rte_sws_3526 rte_sws_3527 rte_sws_3530 rte_sws_3531 rte_sws_3532 rte_sws_2697 rte_sws_3523 rte_sws_3520 rte_sws_3524 rte_sws_2203 rte_sws_1131 rte_sws_2512 rte_sws_1133 rte_sws_1359 rte_sws_1166 rte_sws_1135 rte_sws_1137
[RTE00073] Data items are atomic	rte_sws_4527
[RTE00075] API for accessing static memory sections	rte_sws_1118 rte_sws_1119
[RTE00077] Instantiation of static memory sections	rte_sws_3790 rte_sws_2303 rte_sws_2304 rte_sws_3789 rte_sws_3782 rte_sws_2305 rte_sws_5062

[RTE00078] Support for INVALIDATE attribute	rte_sws_5024 rte_sws_2594 rte_sws_2702 rte_sws_1206 rte_sws_1282 rte_sws_1231 rte_sws_5063 rte_sws_2626 rte_sws_3800 rte_sws_3801 rte_sws_3802 rte_sws_5064 rte_sws_3778 rte_sws_2599 rte_sws_2600 rte_sws_2603 rte_sws_2629 rte_sws_2607 rte_sws_2666 rte_sws_2589 rte_sws_2590 rte_sws_2609
[RTE00079] Single asynchronous client-server interaction	rte_sws_3765 rte_sws_3766 rte_sws_3771 rte_sws_3772 rte_sws_2658 rte_sws_1105 rte_sws_1109 rte_sws_1133 rte_sws_1359 rte_sws_1166
[RTE00080] Multiple requests of servers	rte_sws_4516 rte_sws_4520
[RTE00082] Standardized communication protocol	rte_sws_2649 rte_sws_2651 rte_sws_2652 rte_sws_2653 rte_sws_2579 rte_sws_5066 rte_sws_2654 rte_sws_2655 rte_sws_2656 rte_sws_2657 rte_sws_5067 rte_sws_5054 rte_sws_5055 rte_sws_6028 rte_sws_5056 rte_sws_5057 rte_sws_5058 rte_sws_5059
[RTE00083] Optimization for source-code components	rte_sws_1274 rte_sws_1152
[RTE00084] Support infrastructural errors	rte_sws_2593 rte_sws_1318
[RTE00087] Application Header File	rte_sws_1000 rte_sws_3786 rte_sws_1004 rte_sws_1006 rte_sws_1263 rte_sws_5078 rte_sws_1009 rte_sws_1132
[RTE00089] Independent access to interface elements	rte_sws_6008
[RTE00091] Inter-ECU Marshalling	rte_sws_4505 rte_sws_4506 rte_sws_4507 rte_sws_4508 rte_sws_2557 rte_sws_5081 rte_sws_4527
[RTE00092] Implementation of VFB model waitpoints	rte_sws_1358 rte_sws_3010 rte_sws_7007 rte_sws_3018
[RTE00094] Communication and Resource Errors	rte_sws_2524 rte_sws_2525 rte_sws_1318 rte_sws_2571 rte_sws_1034 rte_sws_1073 rte_sws_1074 rte_sws_2674 rte_sws_1207 rte_sws_1339 rte_sws_1084 rte_sws_3774 rte_sws_1086 rte_sws_1093 rte_sws_2598 rte_sws_1094 rte_sws_1095 rte_sws_2572 rte_sws_1103 rte_sws_1104 rte_sws_1105 rte_sws_1106 rte_sws_1107 rte_sws_1112 rte_sws_1113 rte_sws_1114 rte_sws_3606 rte_sws_2578 rte_sws_3803 rte_sws_2602 rte_sws_1261 rte_sws_1262 rte_sws_1259 rte_sws_1260
[RTE00098] Explicit Transmission	rte_sws_6011 rte_sws_6016 rte_sws_1071
[RTE00099] Decoupling of interrupts	rte_sws_3600 rte_sws_3594 rte_sws_3530 rte_sws_3531 rte_sws_3532

[RTE00100] Compiler independent API	rte_sws_1314
[RTE00107] Support for INFORMATION_TYPE attribute	rte_sws_6010 rte_sws_4500 rte_sws_2516 rte_sws_2518 rte_sws_2520 rte_sws_2521 rte_sws_2522 rte_sws_2523 rte_sws_2524 rte_sws_2525 rte_sws_2571 rte_sws_2572 rte_sws_1135 rte_sws_1137
[RTE00108] Support for INIT_VALUE attribute	rte_sws_4525 rte_sws_6009 rte_sws_4501 rte_sws_4502 rte_sws_2517 rte_sws_1268 rte_sws_5078
[RTE00109] Support for RECEIVE_MODE attribute	rte_sws_3018 rte_sws_6002 rte_sws_6012 rte_sws_2519
[RTE00110] Support for BUFFERING attribute	rte_sws_2515 rte_sws_2522 rte_sws_2523 rte_sws_2524 rte_sws_2525 rte_sws_2526 rte_sws_2527 rte_sws_2529 rte_sws_2530 rte_sws_7008 rte_sws_2571 rte_sws_2572
[RTE00111] Support for CLIENT_MODE attribute	rte_sws_1293 rte_sws_1294
[RTE00115] API for data consistency mechanism	rte_sws_1120 rte_sws_1307 rte_sws_1122 rte_sws_1308
[RTE00116] RTE Initialization, finalization and resumption	rte_sws_2513 rte_sws_2535 rte_sws_2536 rte_sws_2538 rte_sws_2544 rte_sws_2569 rte_sws_2585 rte_sws_2570 rte_sws_2584
[RTE00121] Support for FILTER attribute	rte_sws_5503 rte_sws_5500 rte_sws_5501
[RTE00122] Support for SUCCESS attribute	rte_sws_5504 rte_sws_3754 rte_sws_3756 rte_sws_3757 rte_sws_3604 rte_sws_3758 rte_sws_1080 rte_sws_2673 rte_sws_1083 rte_sws_1283 rte_sws_2676 rte_sws_1284 rte_sws_2677 rte_sws_1285 rte_sws_2678 rte_sws_1286 rte_sws_2679 rte_sws_1287 rte_sws_1084 rte_sws_1086 rte_sws_3002 rte_sws_3775 rte_sws_2612 rte_sws_2610 rte_sws_3005 rte_sws_3776 rte_sws_5065 rte_sws_2611 rte_sws_1137
[RTE00123] Forwarding of application level errors	rte_sws_2593 rte_sws_2576 rte_sws_1103 rte_sws_2577 rte_sws_2578
[RTE00124] APIs for application level server errors	rte_sws_2573 rte_sws_2575 rte_sws_1103 rte_sws_1130
[RTE00125] Interaction of 1:n communication with the SUCCESS attribute	rte_sws_5506
[RTE00126] C support	rte_sws_3724 rte_sws_1005 rte_sws_3709 rte_sws_3710 rte_sws_1162 rte_sws_1169 rte_sws_1167
[RTE00128] Implicit Reception	rte_sws_3012 rte_sws_6000 rte_sws_6001 rte_sws_6004 rte_sws_6011 rte_sws_3741 rte_sws_1268 rte_sws_1005 rte_sws_3709 rte_sws_3710

[RTE00129] Implicit Transmission	rte_sws_6011 rte_sws_3570 rte_sws_3571 rte_sws_3572 rte_sws_3573 rte_sws_3744 rte_sws_3746 rte_sws_5509
[RTE00130] API to determine executing runnable entity	protection is cancelled for release 3.0
[RTE00131] n:1 Sender-receiver communication	rte_sws_2670 rte_sws_3760 rte_sws_3761 rte_sws_3762 rte_sws_1071 rte_sws_1072 rte_sws_1077 rte_sws_1081 rte_sws_2633 rte_sws_2635 rte_sws_2631 rte_sws_1091 rte_sws_1092 rte_sws_1135
[RTE00133] No parallel execution of runnable instance	rte_sws_2697 rte_sws_2698 rte_sws_3523
[RTE00134] Runnable entity categories supported by the RTE	rte_sws_3016 rte_sws_6003 rte_sws_6007 rte_sws_3574 rte_sws_3954 rte_sws_3598 rte_sws_3955 rte_sws_3599 rte_sws_3953 rte_sws_3956 rte_sws_3957
[RTE00137] API for mismatched ports	rte_sws_1368 rte_sws_1369
[RTE00138] C++ support	rte_sws_1370 rte_sws_3724 rte_sws_1162 rte_sws_1169 rte_sws_1011
[RTE00139] API for unconnected ports	rte_sws_3019 rte_sws_2750 rte_sws_3978 rte_sws_5101 rte_sws_3980 rte_sws_5102 rte_sws_1329 rte_sws_5100 rte_sws_1330 rte_sws_1331 rte_sws_1336 rte_sws_1344 rte_sws_1345 rte_sws_1332 rte_sws_3783 rte_sws_1346 rte_sws_1347 rte_sws_3784 rte_sws_3785 rte_sws_2638 rte_sws_2639 rte_sws_2640 rte_sws_2641 rte_sws_2642 rte_sws_1333 rte_sws_1337 rte_sws_1334 rte_sws_5099
[RTE00140] Binary-code AUTOSAR software components	rte_sws_1315 rte_sws_1000 rte_sws_1195
[RTE00141] Explicit Reception	rte_sws_6011 rte_sws_1072 rte_sws_1091 rte_sws_1092
[RTE00142] InterRunnableVariables	rte_sws_3518 rte_sws_3588 rte_sws_3591 rte_sws_3589 rte_sws_3516 rte_sws_3517 rte_sws_3582 rte_sws_3583 rte_sws_3584 rte_sws_3519 rte_sws_3580 rte_sws_3550 rte_sws_1303 rte_sws_3581 rte_sws_3552 rte_sws_3556 rte_sws_3558 rte_sws_3553 rte_sws_1304 rte_sws_3557 rte_sws_3559 rte_sws_3555 rte_sws_3560 rte_sws_1305 rte_sws_3562 rte_sws_3563 rte_sws_3564 rte_sws_3565 rte_sws_1306 rte_sws_3567 rte_sws_3568 rte_sws_3569 rte_sws_2636 rte_sws_1350 rte_sws_1351

[RTE00143] Mode switches	rte_sws_2706 rte_sws_2500 rte_sws_2662 rte_sws_2663 rte_sws_2664 rte_sws_2503 rte_sws_2504 rte_sws_2667 rte_sws_2661 rte_sws_2562 rte_sws_2564 rte_sws_2563 rte_sws_2587 rte_sws_2665 rte_sws_2668 rte_sws_2544 rte_sws_2630 rte_sws_2669 rte_sws_2546 rte_sws_2634 rte_sws_2631 rte_sws_2675 rte_sws_2512
[RTE00144] Mode switch notification via AUTOSAR interfaces	rte_sws_2544 rte_sws_2549 rte_sws_2586 rte_sws_2508 rte_sws_2566 rte_sws_2624 rte_sws_2567 rte_sws_2546 rte_sws_2627 rte_sws_2659 rte_sws_2568 rte_sws_2628 rte_sws_2660
[RTE00145] Compatibility mode	rte_sws_1257 rte_sws_3794 rte_sws_1279 rte_sws_1326 rte_sws_1277 rte_sws_1151 rte_sws_1216 rte_sws_1234
[RTE00146] Vendor mode	rte_sws_1234
[RTE00147] Support for communication infrastructure time-out notification	rte_sws_5020 rte_sws_5021 rte_sws_3759 rte_sws_5022 rte_sws_2703 rte_sws_2599 rte_sws_2600 rte_sws_2604 rte_sws_2629 rte_sws_2607 rte_sws_2666 rte_sws_2589 rte_sws_2590 rte_sws_2609
[RTE00148] Support 'Specification of Memory Mapping'	rte_sws_3788
[RTE00149] Support 'Specification of Compiler Abstraction'	rte_sws_3787 rte_sws_1164
[RTE00150] Support 'Specification of Platform Types'	rte_sws_1164
[RTE00151] Support RTE relevant requirements of the 'General Requirements on Basic Software Modules'	see [BSW...] entries in this table
[RTE00152] Support for port-defined argument values	rte_sws_1360
[RTE00153] Support of Measurement	rte_sws_3951 rte_sws_3900 rte_sws_3972 rte_sws_3973 rte_sws_3974 rte_sws_3901 rte_sws_3975 rte_sws_3976 rte_sws_3977 rte_sws_3902 rte_sws_3978 rte_sws_5101 rte_sws_3980 rte_sws_5102 rte_sws_3979 rte_sws_3903 rte_sws_3904 rte_sws_3950 rte_sws_3981 rte_sws_3982

[RTE00154] Support of Calibration	rte_sws_3970 rte_sws_3958 rte_sws_7186 rte_sws_3959 rte_sws_7185 rte_sws_3905 rte_sws_3906 rte_sws_3907 rte_sws_3971 rte_sws_3909 rte_sws_3942 rte_sws_3910 rte_sws_3943 rte_sws_3911 rte_sws_3912 rte_sws_3968 rte_sws_3913 rte_sws_3947 rte_sws_3936 rte_sws_3914 rte_sws_3948 rte_sws_3915 rte_sws_3935 rte_sws_3916 rte_sws_3917 rte_sws_3918 rte_sws_3969 rte_sws_3908 rte_sws_3920 rte_sws_3940 rte_sws_3921 rte_sws_3922 rte_sws_3960 rte_sws_3932 rte_sws_3933 rte_sws_3934 rte_sws_3961 rte_sws_3962 rte_sws_3963 rte_sws_3964 rte_sws_3965 rte_sws_3966 rte_sws_3967 rte_sws_3937 rte_sws_3938 rte_sws_6029 rte_sws_3949
[RTE00155] API to access calibration parameters	rte_sws_1252 rte_sws_1300 rte_sws_1254 rte_sws_3927 rte_sws_3952 rte_sws_3928 rte_sws_3929 rte_sws_3930 rte_sws_6029 rte_sws_3949
[RTE00156] Support different calibration data emulation methods	rte_sws_3970 rte_sws_3905 rte_sws_3906 rte_sws_3971 rte_sws_3909 rte_sws_3942 rte_sws_3910 rte_sws_3943 rte_sws_3911 rte_sws_3968 rte_sws_3913 rte_sws_3947 rte_sws_3936 rte_sws_3914 rte_sws_3948 rte_sws_3915 rte_sws_3935 rte_sws_3916 rte_sws_3917 rte_sws_3918 rte_sws_3969 rte_sws_3908 rte_sws_3920 rte_sws_3940 rte_sws_3922 rte_sws_3960 rte_sws_3932 rte_sws_3933 rte_sws_3934 rte_sws_3961 rte_sws_3962 rte_sws_3963 rte_sws_3964 rte_sws_3965 rte_sws_3966 rte_sws_3967
[RTE00157] Support calibration parameters in NVRAM	rte_sws_3936 rte_sws_3937 rte_sws_3938
[RTE00158] Support separation of calibration parameters	rte_sws_3907 rte_sws_3911 rte_sws_3912 rte_sws_3908
[RTE00159] Sharing of calibration parameters	rte_sws_2750 rte_sws_3958 rte_sws_7186 rte_sws_7185
[RTE00160] Debounced start of runnable entities	rte_sws_2697
[RTE00161] Activation Offset of runnable entities	rte_sws_7000

2 RTE Overview

2.1 The RTE in the Context of AUTOSAR

The Run-Time Environment (RTE) is at the heart of the AUTOSAR ECU architecture. The RTE is the realization (for a particular ECU) of the interfaces of the AUTOSAR Virtual Function Bus (VFB). The RTE provides the infrastructure services that enable communication to occur between AUTOSAR software-components as well as acting as the means by which AUTOSAR software-components access basic software modules including the OS and communication service.

The RTE encompasses both the variable elements of the system infrastructure that arise from the different mappings of components to ECUs as well as standardized RTE services.

The RTE is generated¹ for each ECU to ensure that the RTE is optimal for the ECU [RTE00023].

2.2 AUTOSAR Concepts

This section introduces fundamental AUTOSAR concepts and how they are understood within the context of the RTE.

2.2.1 AUTOSAR Software-components

In AUTOSAR, “application” software is conceptually located above the AUTOSAR RTE and consists of “AUTOSAR application software-components” that are ECU and location independent and “AUTOSAR sensor-actuator components” that are dependent on ECU hardware and thus not readily relocatable for reasons of performance/efficiency. This means that, subject to constraints imposed by the system designer, an AUTOSAR software-component can be deployed to any available ECU during system configuration. The RTE is then responsible for ensuring that components can communicate and that the system continues to function as expected wherever the components are deployed. Considering sensor/actuator software components, they may only directly address the local ECU abstraction. Therefore, access to remote ECU abstraction shall be done through an intermediate sensor/actuator software component which broadcasts the information on the remote ECU. Hence, moving the sensor/actuator software components on different ECUs, may then imply to also move connected devices (sensor/actuator) to the same ECU (provided that efficient access is needed).

¹An implementation is free to *configure* rather than *generate* the RTE. The remainder of this specification refers to generation for reasons of simplicity only and these references should not be interpreted as ruling out either a wholly configured, or partially generated and partially configured, RTE implementation.

An AUTOSAR software-component is defined by a *type* definition that defines the component's interfaces. A component type is instantiated when the component is deployed to an ECU. A component type can be instantiated more than once on the same ECU in which case the component type is said to be "multiply instantiated". The RTE supports per-instance memory sections that enable each component instance to have private states.

The RTE supports both AUTOSAR software-components where the source is available ("source-code software-components") [RTE00024] and AUTOSAR software-components where only the object code ("object-code software components") is available [RTE00140].

Details of AUTOSAR software-components in relation to the RTE are presented in Section 4.1.2.

2.2.2 Basic Software Modules

As well as "AUTOSAR software-components" an AUTOSAR ECU includes basic software modules. Basic software modules can access the ECU abstraction layer as well as other basic software modules directly and are thus neither ECU nor location independent.

An "AUTOSAR software-component" *cannot* directly access basic software modules – all communication is via AUTOSAR interfaces and therefore under the control of the RTE. The requirement to not have direct access applies to all basic software modules including the operating system [RTE00020] and the communication service.

2.2.3 Communication

The communication interface of an AUTOSAR software-component consists of several ports (which are characterized by port-interfaces). An AUTOSAR software-component can communicate through its interfaces with other AUTOSAR software-components (whether that component is located on the same ECU or on a different ECU) or with basic software modules that have a port and are located on the same ECU. This communication can *only* occur via the component's ports. A port can be categorized by either a sender-receiver or client-server port-interface. A sender-receiver interface provides a message passing facility whereas a client-server interface provides function invocation.

2.2.3.1 Communication Models

The AUTOSAR VFB Specification [13] defines two communication models within the RTE core services; sender-receiver (signal passing) and client-server (function invocation). Each communication model can be applied to either intra-ECU software-

component distribution (which includes both intra-task and inter-task distribution) and inter-ECU software-component distribution. Intra-task communication occurs between runnable entities that are mapped to the same OS task whereas inter-task communication occurs between runnable entities mapped to different tasks and can therefore involve a context switch and possibly cross memory protection boundaries. In contrast, inter-ECU communication occurs between runnable entities in components that have been mapped to different ECUs and so is inherently concurrent and involves potentially unreliable communication.

Details of the communication models that are supported by the RTE are contained in Section 4.3.

2.2.3.2 Communication Modes

The RTE supports two modes for sender-receiver communication:

- **Explicit** — A component uses explicit RTE API calls to send and receive data elements [RTE00098].
- **Implicit** — The RTE automatically reads a specified set of data elements before a runnable is invoked and automatically writes (a different) set of data elements after the runnable entity has terminated [RTE00128] [RTE00129]. The term “implicit” is used here since the runnable does not actively initiate the reception or transmission of data.

Implicit and explicit communication is considered in greater detail in Section 4.3.1.5.

2.2.3.3 Static Communication

[rte_sws_6026] The RTE shall support static communication only.

Static communication includes only those communication connections where the source(s) and destination(s) of all communication is known at the point the RTE is generated. [RTE00025]. Dynamic reconfiguration of communication is not supported due to the run-time and code overhead which would therefore limit the range of devices for which the RTE is suitable.

2.2.3.4 Multiplicity

As well as point to point communication (i.e. “1:1”) the RTE supports communication connections with multiple providers or requirers:

- When using sender-receiver communication, the RTE supports both “1:n” (single sender with multiple receivers) [RTE00028] and “n:1” (multiple senders and a single receiver) [RTE00131] communication with the restriction that multiple

senders are not allowed for `mode switch notifications`, see metamodel restrictions `rte_sws_2670`.

The execution of the multiple senders or receivers is not coordinated by the RTE. This means that the actions of different software-components are independent – the RTE does not ensure that different senders transmit data simultaneously and does not ensure that all receivers read data or receive events simultaneously.

- When using client-server communication, the RTE supports “n:1” (multiple clients and a single server) [RTE00029] communication. The RTE does *not* support “1:n” (single client with multiple servers) client-server communication.

Irrespective of whether “1:1”, “n:1” or “1:n” communication is used, the RTE is responsible for implementing the communication connections and therefore the AUTOSAR software-component is unaware of the configuration. This permits an AUTOSAR software-component to be redeployed in a different configuration without modification.

2.2.4 Concurrency

AUTOSAR software-components have no direct access to the OS and hence there are no “tasks” in an AUTOSAR application. Instead, concurrent activity within AUTOSAR is based around *runnable entities* within components that are invoked by the RTE.

The AUTOSAR VFB specification [13] defines a runnable entity as a “sequence of instructions that can be started by the Run-Time Environment”. A component provides one² or more runnable entities [RTE00031] and each runnable entity has exactly one entry point. An entry point defines the *symbol* within the software-component’s code that provides the implementation of a runnable entity.

The RTE is responsible for invoking runnable entities – AUTOSAR software-components are not able to (dynamically) create private threads of control. Hence, all activity within an AUTOSAR application is initiated by the triggering of runnable entities by the RTE as a result of `RTEEvents`.

An *RTEEvent* encompasses all possible situations that can trigger execution of a runnable entity by the RTE. The different classes of `RTEEvent` are defined in Section 5.7.5.

The RTE supports runnable entities in any component that has an AUTOSAR interface - this includes AUTOSAR software-components and basic software modules.³

Runnable entities are divided into multiple categories with each category supporting different facilities. The categories supported by the RTE are described in Section 4.2.2.2.

²The VFB specification does not permit zero runnable entities.

³The OS and COM are basic software modules but present a *standardized interface* to the RTE and have no AUTOSAR interface. The OS and COM therefore do not have runnable entities.

2.3 The RTE Generator

The RTE generator is one of a set of tools⁴ that create the realization of the AUTOSAR virtual function bus for an ECU based on information in the *ECU Configuration Description*. The RTE Generator is responsible for creating the AUTOSAR software-component API functions that link AUTOSAR software-components to the OS and manage communication between AUTOSAR software-components and between AUTOSAR software-components and basic software modules.

The RTE generation process consists of two distinct phases:

- **RTE Contract phase** – a limited set of information about a component, principally the AUTOSAR interface definitions, is used to create an application header file for a component type. The application header file defines the “contract” between component and RTE.
- **RTE Generation phase** - all relevant information about components, their deployment to ECUs and communication connections is used to generate the RTE. One RTE is generated for each ECU in the system.

The two-phase development model ensures that the RTE generated application header files are available for use for source-code AUTOSAR software-components as well as object-code AUTOSAR software-components with both types of component having access to all definitions created as part of the RTE generation process.

The RTE generation process, and the necessary inputs in each phase, are considered in more detail in Section 3.

2.4 Design Decisions

This section details decisions that affect both the general direction that has been taken as well as the actual content of this document.

1. The role of this document is to specify RTE behavior, not RTE implementation. Implementation details should not be considered to be part of the RTE software specification unless they are explicitly marked as RTE requirements.
2. An AUTOSAR system consists of multiple ECUs each of which contains an RTE that may have been generated by different RTE generators. Consequently, the specification of how RTEs from multiple vendors interoperate is considered to be within the scope of this document.
3. The RTE does not have sufficient information to be able to derive a mapping from runnable entity to OS task. The decision was therefore taken to require that the mapping be specified as part of the RTE input.

⁴The RTE generator works in conjunction with other tools, for example, the OS and COM generators, to fully realize the AUTOSAR VFB.

4. Support for C++ is provided by making the C RTE API available for C++ components rather than specifying a completely separate object-oriented API. This decision was taken for two reasons; firstly the same interface for the C and C++ simplifies the learning curve and secondly a single interface greatly simplifies both the specification and any subsequent implementations.
5. There is no support within the specification for Java.
6. The support for AUTOSAR OS protection mechanisms has been deferred until a later release of the RTE software specification.
7. The AUTOSAR meta-model is a highly expressive language for defining systems however for reasons of practicality certain restrictions and constraints have been placed on the use of the meta-model. The restrictions are described in Appendix A.

3 RTE Generation Process

This chapter describes the methodology of the RTE generation. For a detailed description of the overall AUTOSAR methodology refer to methodology document [2].

[rte_sws_2514] The RTE generator shall produce the same RTE API and RTE code when the input information is the same.

The RTE-Generator gets involved in the AUTOSAR Methodology twice. In the following section the two applications of the RTE-Generator are described.

In Figure 3.1 the overall AUTOSAR Methodology is outlined with respect to the RTE.

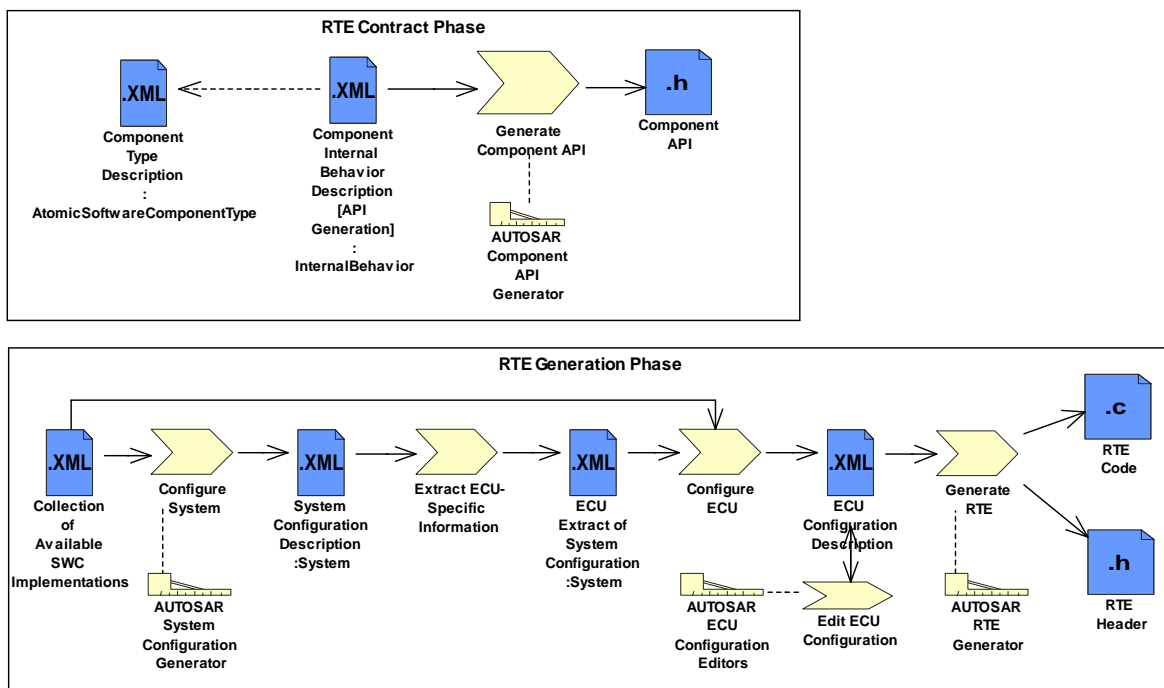


Figure 3.1: System Build Methodology

For the development of AUTOSAR Software Components it is essential that the 'Component API Generator Tool' ¹ produces the 'Component API' file in the so called 'RTE Contract Phase' (see section 3.1).

The whole vehicle functionality is described with means of Composite SW-Components and Atomic SW-Components. In the Composite SW-Component descriptions the connections between the SW-Component's ports are also defined. Such a collection of SW-Components connected to each other, without the mapping on actual ECUs, is called the VFB view.

¹The 'Component API Generator Tool' might be a separate tool or the RTE-Generator might be operated in a special mode to achieve the same functionality. This specification does not require how the tool is implemented.

During the 'Configure System' step the 'System Configuration Generator' gets the information about the needed SW-Components, the available ECUs and the System Constraints. Now the Atomic SW-Components are mapped on the available ECUs.

Since in the VFB view the communication relationships between the Atomic SW-Components have been described and the mapping of each Atomic SW-Component to a specific ECU has been fixed, the communication matrix can be generated. In the SW-Component descriptions the signals that are exchanged through ports are defined in an abstract way. Now the 'System Configuration Generator' needs to define system signals (including the actual signal length and the frames in which they will be transmitted) to be able to transmit the application signals over some network. COM signals that correspond to the system signals will be later used by the 'RTE Generator' to actually transmit the application signals.

In the next step the 'System Configuration Description' is split into descriptions for each individual ECU. The extract only contains information necessary to configure each ECU individually and it is fed into the ECU Configuration for each ECU.

[rte_sws_5000] The RTE is configured and generated for each ECU instance individually.

The 'ECU Configuration Editors' (see also Section 3.2) are working iteratively on the 'ECU Configuration Description' until all configuration issues are resolved. There will be the need for several configuration editors, each specialized on a specific part of ECU Configuration. So one editor might be configuring the COM stack (not the communication matrix but the interaction of the individual modules) while another editor is used to configure the RTE.

Since the configuration of a specific Basic-SW module is not entirely independent from other modules there is the need to apply the editors several times to the 'ECU Configuration Description' to ensure all configuration parameters are consistent.

Only when the configuration issues are resolved the 'RTE Generator' will be used to generate the actual RTE code (see also Section 3.3) which will then be compiled and linked together with the other Basic-SW modules and the SW-Components code.

The 'RTE Generator' needs to cope with many sources of information since the necessary information for the RTE Generator is based on the 'ECU Configuration Description' which might be distributed over several files and itself references to multiple other AUTOSAR descriptions.

[rte_sws_5001] The RTE Generation tools needs to support input according to the Interoperability of AUTOSAR Authoring Tools document [9].

This is just a rough sketch of the main steps necessary to build an ECU with AUTOSAR and how the RTE is involved in this methodology. For a more detailed description of the AUTOSAR Methodology please refer to the methodology document [2]. In the next sections the steps with RTE interaction are explained in more detail.

3.1 RTE Contract Phase

To be able to support the SW-Component development with RTE-specific APIs the 'Component API' (application header file) is generated from the 'SW-Component Internal Behavior Description' (see Figure 3.1) by the RTE-Generator in the so called 'RTE Contract Phase' (see Figure 3.2).

In the SW-Component Interface description – which is using the AUTOSAR Software Component Template – at least the AUTOSAR interfaces of the particular SW-Component have to be described. This means the SW-Component Types with Ports and their Interfaces. In the SW-Component Internal Behavior description additionally the Runnable Entities and the RTE Events are defined. From this information the RTE-Generator can generate specific APIs to access the Ports and send and receive data.

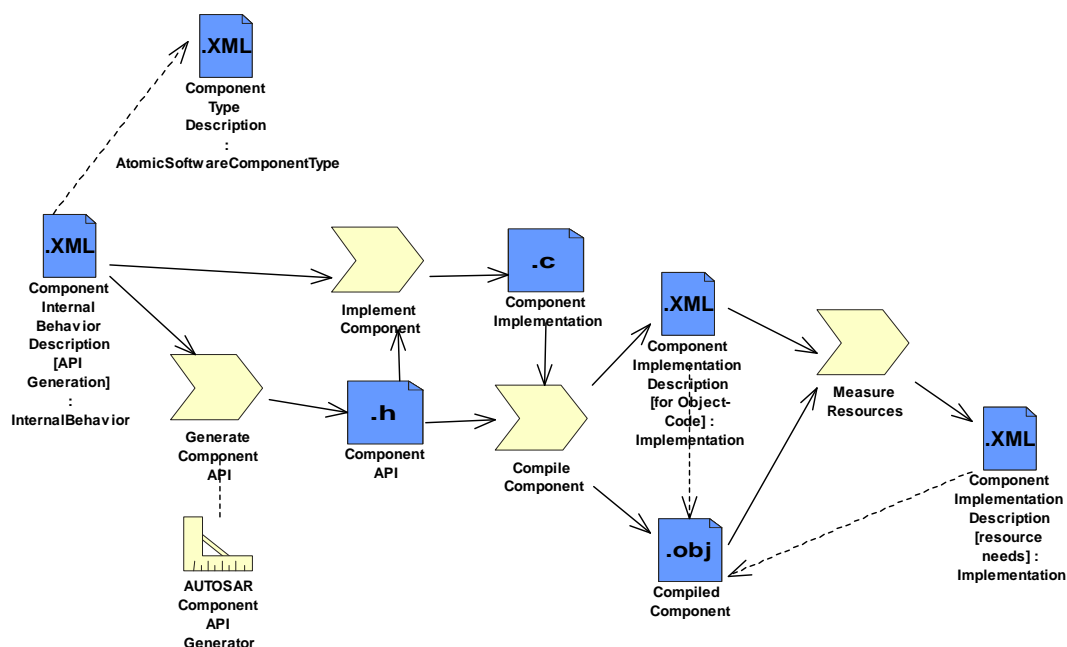


Figure 3.2: RTE Contract Phase

With the generated 'Component API' (application header file) the Software Component developer can provide the Software Component's source code without being concerned as to whether the communication will later be local or using some network(s).

It has to be considered that the SW-Component development process is iterative and that the SW-Component description might be changed during the development of the SW-Component. This requires the application header file to be regenerated to reflect the changes done in the SW-Component description.

When the SW-Component has been compiled successfully the 'Component Implementation Description Generation' tool will analyze the resulting object files and enhance the SW-Component description with the information from the specific implementation. This includes information about the actual memory needs for ROM as well as for

RAM and goes into the 'Component Implementation Description' section of the SW-Component Description template.

So when a SW-Component is delivered it will consist of the following parts:

- Component Type Description
- Component Internal Behavior Description
- The actual source and/or object code
- Component Implementation Description

The afore listed information will be needed to provide enough information for the System Generation steps when the whole system is assembled.

3.2 RTE Configuration Editing

During the configuration of an ECU the RTE also needs to be configured. This is mainly divided into two sections: The configuration of the RTE and the request for configuration of other modules.

So first the 'RTE Configuration Editor' needs to collect all the information needed to establish an operational RTE. This gathering includes information on the SW-Component instances and their communication relationships, the Runnable Entities and the involved RTE-Events and so on. The main source for all this information is the 'ECU Configuration Description', which might provide references to further descriptions like the SW-Component description or the System Configuration description.

One extremely important point is the mapping of application signals from SW-Component's ports to COM signals. A mapping of the application signals to system signals has already been defined by the 'System Configuration Generator' (see Figure 3.1). The 'RTE Configuration Editor' now has to substantiate this system-level mapping by mapping the application signals to COM signals for the ECU. This application signal to COM signal mapping has to respect the mapping from application signals to system signals done at system generation time.

Additional requirements on the interaction between RTE and OS are specified in section 3.4.

The usage of 'ECU Configuration Editors' covering different parts of the 'ECU Configuration Description' will – if there are no cyclic dependencies which do not converge – converges to a stable configuration and then the ECU Configuration process is finished. A detailed description of the ECU Configuration can be found in [7]. The next phase is the generation of the actual RTE.

3.3 RTE Generation Phase

After the ECU has been entirely configured the generation of the actual RTE can be performed. Since all the relationships to and from the other Basic-SW modules have been already resolved during the ECU Configuration phase, the generation can be performed in parallel for all modules (see Figure 3.3).




Figure 3.3: RTE Generation Phase

The actual SW-Component and Basic-SW modules code will be linked together with the RTE code to build the entire ECU software.

3.4 RTE Configuration and OS Interaction

The generated RTE interacts with AUTOSAR COM and OS. For the latter, the RTE both uses OS objects already in existence (e.g. tasks for which the RTE generator builds bodies) as well as requires new objects (e.g. a schedule table or periodic alarms for periodic runnable entities). The coordination of configuration information between the OS and RTE is therefore key since both the RTE and OS have to agree upon the set of OS objects.

The AUTOSAR OS is configured in the ECU Configuration Description² [7]. The RTE configurator/generator needs to communicate its needs to the OS and therefore it seems sensible to use the same format order to allow the communication of the set of OS object required by the generated RTE.

The specification of the OS objects used by the generated RTE, henceforth termed *OsNeeds*, can be done either at configuration time only or at a mixture of configuration and generation time, depending on which approach is supported by the configuration and generation tools of RTE and OS. Thus according to figure 3.4 the output information *OsNeeds* can be alternatively provided by the RTE Configuration Editor or the RTE Generator.

[rte_sws_5071] If the RTE Generator (in the generation phase) does not export *OsNeeds* information then the RTE Configuration editor shall export the *OsNeeds* information.

²This is the same format which is also used to store the configuration values for the RTE and the OS.

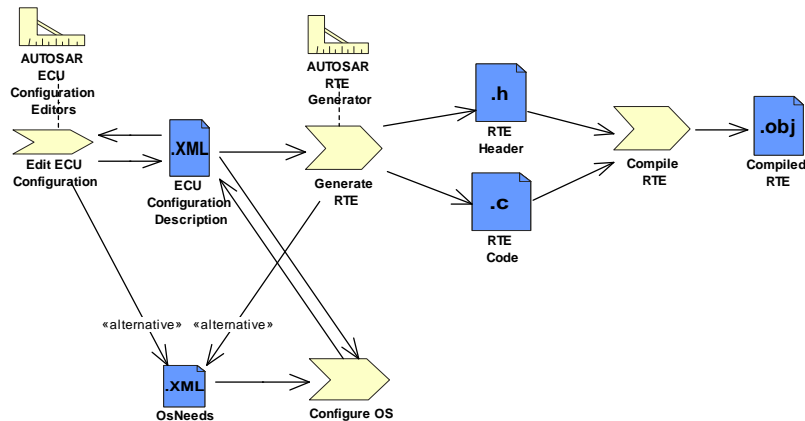


Figure 3.4: RTE Editing Phase and OS interaction

[rte_sws_5076] When provided with `OsNeeds` information, the RTE Generator (in the generation phase) shall utilize only the OS objects defined in the container.

Requirement 5076 indicates that it is not possible to use/introduce new OS objects which are not part of the `OsNeeds` definition when such objects are provided.

[rte_sws_5077] The RTE shall access `OsNeeds` information from the module configuration container `Rte_OsNeeds`.

Requirement 5077 **names** the container that holds the `OsNeeds` information; and hence the container that the RTE generator should read and write. This container is separate from the OS information that is purely read by the RTE configurator/generator (e.g. tasks) and the two containers will require merging within the OS generation process.

If the RTE Configuration editor provides a list of the OS Objects to be used the RTE Generator then that information does not need to be recreated by the RTE generator during generation phase and hence the generation phase should respect the objects provided by the `OsNeeds` information.

[rte_sws_5072] If no `Rte_OsNeeds` container is provided, the RTE Generator (in the generation phase) shall assume the existence of required OS objects *ex nihilo* and export `OsNeeds` information containing all required OS objects.

If RTE Configuration editor is not able to provide the `OsNeeds` information (for example, because a Generic ECU Configuration Editor is used that does not know the details of a particular RTE implementation) the RTE Generator in the generation phase shall provide this information.

[rte_sws_5073] The exported `OsNeeds` shall be described using the OS Configuration format defined in the OS SWS [11] and formalized according to the ECU Configuration Specification [7].

The exported `OsNeeds` information is, by definition, the sole mechanism for the RTE to impose existence requirements on the OS configuration process. The exported information should contain all information necessary for the OS configuration to proceed.

Since the OS Configuration format is used to describe the `OsNeeds` this information can be directly used within the OS Configuration process to enhance the information.

Caution has to be taken in order not to break the references between the RTE Configuration and the `OsNeeds`. If changes on the `OsNeeds` are performed, for example a task priority changed, and this requires changes on the RTE configuration/generation this has to be synchronized.

[rte_sws_5074] The exported `OsNeeds` shall only utilize the standardized configuration parameters from the OS SWS [11] in an ICC3 delivery.

Because only the standardized configuration parameters are used the RTE and OS are independent and can be provided by different vendors. If both RTE and OS are provided by the same vendor and they are delivered in an ISS2 cluster (bundle) requirement `rte_sws_5074` is not valid.

4 RTE Functional Specification

4.1 Architectural concepts

4.1.1 Scope

In this section the concept of an AUTOSAR software-component and its usage within the RTE is introduced.

The Software-Component Template [17] defines the kinds of SW-Components within the AUTOSAR context. These are shown in Figure 4.1. The abstract `ComponentType` can not be instantiated, so there can only be either a `CompositionType` or a specialized class `ApplicationSoftwareComponentType`, `SensorActuatorSoftwareComponentType`, `ServiceComponentType`, `ComplexDeviceDriverComponentType` or `EcuAbstractionComponentType` of the abstract class `AtomicSoftwareComponentType`.

In the following document the term `AtomicSoftwareComponentType` is used as collective term for all the mentioned non-abstract derived meta-classes.

The `ComponentType` is defining the type of a SW-Component which is independent of any usage and can be potentially re-used several times in different scenarios. In a composition the types are occurring in specific roles which are called `ComponentPrototypes`. The prototype is the utilization of a type within a certain scenario. In AUTOSAR any `ComponentType` can be used as a type for a prototype.

The SW-Components shown in Figure 4.1 are located above the RTE in the architectural Figure 4.2.

Below the RTE there are also software entities that have an AUTOSAR Interface. These are the AUTOSAR services, the ECU Abstraction and the Complex Device Drivers. For these software not only the AUTOSAR Interface will be described but also information about their internal structure will be available in the Basic Software Module Description.

In the next sections the different SW-Components kinds will be described in detail with respect to their influence on the RTE.

4.1.2 RTE and AUTOSAR Software-Components

The description of a SW-Component is divided into the sections

- hierarchical structure
- ports and interfaces
- internal behavior
- implementation

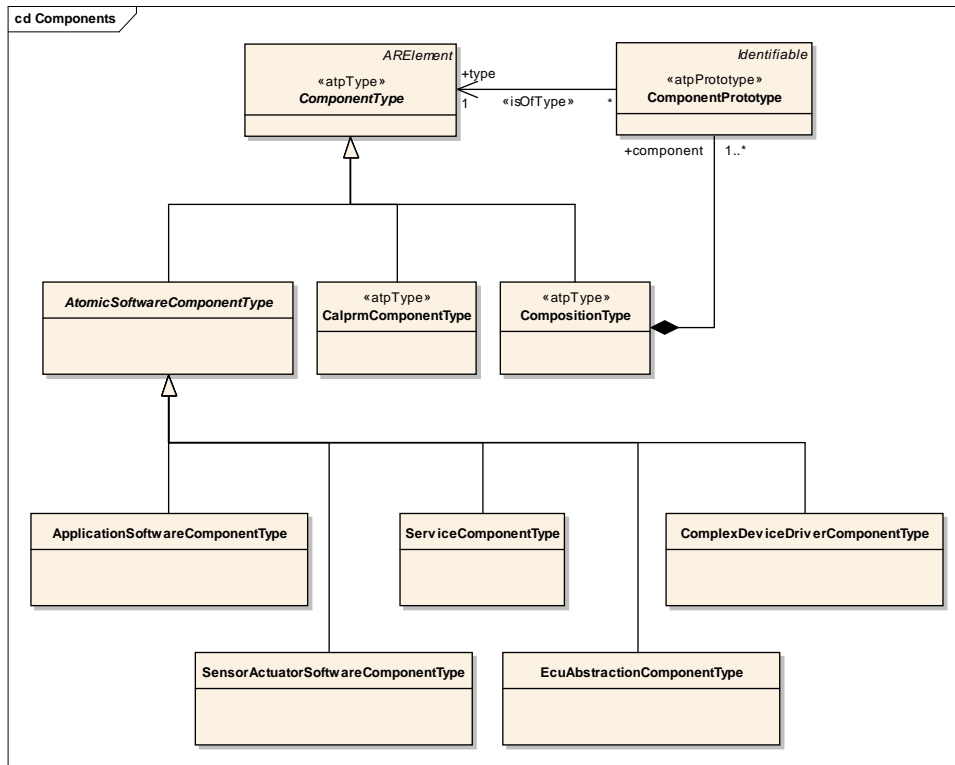


Figure 4.1: AUTOSAR SW-Component classification

which will be addressed separately in the following sections.

4.1.2.1 Structure of SW-Components

In AUTOSAR the structure of an E/E-system is described using the AUTOSAR SW-Component Template and especially the mechanism of compositions. Such a Top Level Composition assembles subsystems and connects their ports.

Of course such a composition utilizes a lot of hierarchical levels where compositions instantiate other composition types and so on. But at some low hierarchical level each composition only consists of AtomicSoftwareComponentType instances. And those instances of AtomicSoftwareComponentTypes are what the RTE is going to be working with.

4.1.2.2 Ports, Interfaces and Connections

Each SW-Component is providing and/or requiring ports to communicate with other SW-Components. This is shown in Figure 4.3. The Interface determines if the port is a sender/receiver or a client/server port.

When compositions are built of instances the ports can be connected either within the composition or made accessible to the outside of the composition. For the connections

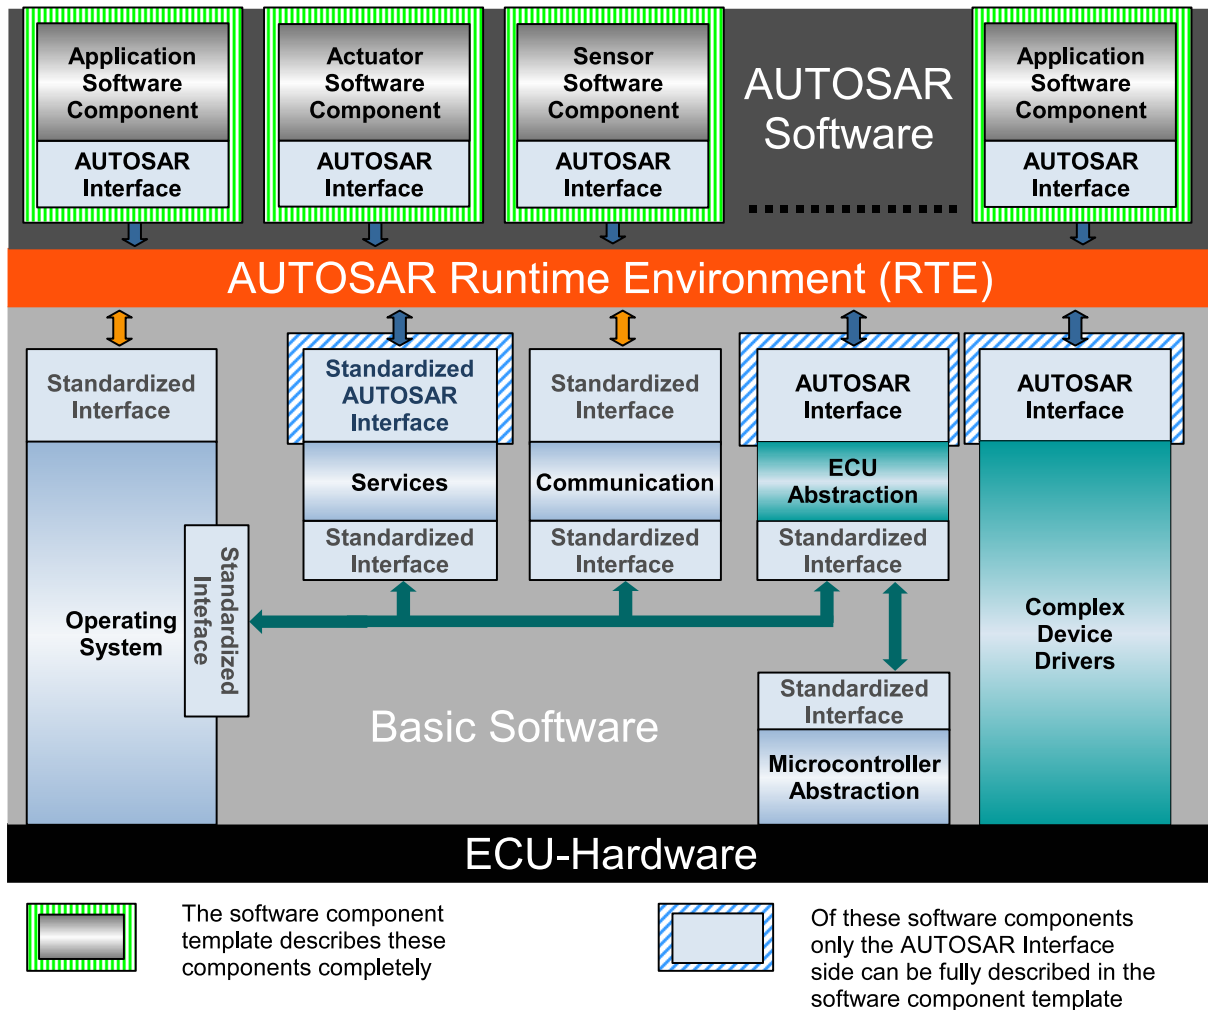


Figure 4.2: AUTOSAR ECU architecture diagram

inside a composition the AssemblyConnector is used, while the DelegationConnector is used to connect ports from the inside of a composition to the outside. Ports not connected will be handled according to the requirement [RTE00139].

The next step is to map the SW-C instances on ECUs and to establish the communication relationships. From this step the actual communication is derived, so it is now fixed if a connection between two instance's ports is going to be over a communication bus or locally within one ECU.

[rte_sws_2200] The RTE shall implement the communication paths specified by the ECU Configuration description (see [RTE00027]).

[rte_sws_2201] The RTE shall implement the semantic of the communication attributes given by the SW-Component description (see [RTE00027]). The semantic of the given communication mechanism shall not change regardless of whether the communication partner is located on the same ECU or remote, the communication is done by COM or the RTE.

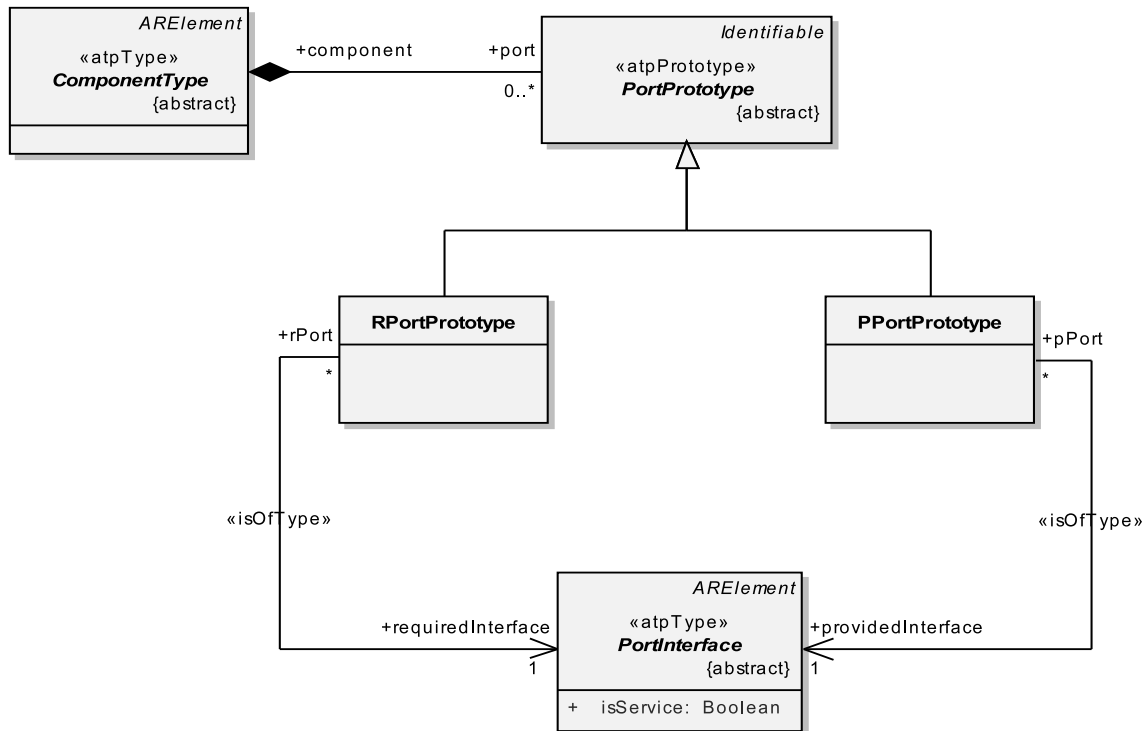


Figure 4.3: SW-Components and Ports

E.g., according to `rte_sws_2200` and `rte_sws_2201` the RTE is not permitted to change the semantic of an asynchronous client to synchronous because both client and server are mapped to the very same ECU.

4.1.2.3 Internal Behavior

Only for AtomicSoftwareComponents the internal structure is exposed in the Internal Behavior description. Here the definition of the Runnable Entities and used RTEEvents is done (see Figure 4.4).

Runnable Entities (also abbreviated simply as Runnable) are the smallest code fragments that are provided by AUTOSAR software-components and those basic software modules that implement AUTOSAR interfaces. They are represented by the meta-class “RunnableEntity”, see Figure 4.5.

In general, software components are composed of multiple Runnable Entities in order to accomplish servers, receivers, feedback, etc.

[rte_sws_2202] The RTE shall support multiple Runnable Entities in AUTOSAR SW-Components (see [RTE00031]).

Runnable Entities are executed in the context of an OS task, their execution is triggered by RTEEvents. Section 4.2.2.2 gives a more detailed description of the concept of Runnable Entities, Section 4.2.2.4 discusses the problem of mapping Runnable Entities

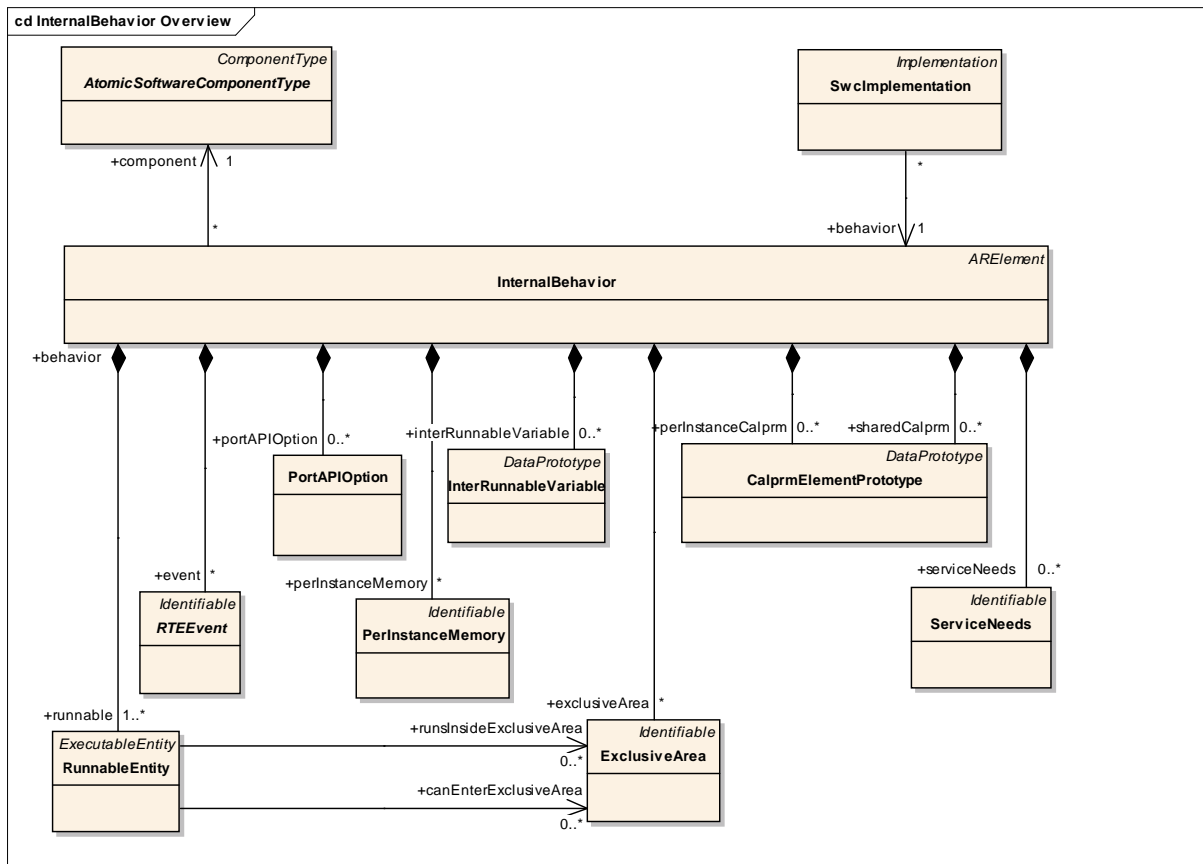


Figure 4.4: SW-Component internal behavior

to OS tasks. RTEEvents and the activation of Runnable Entities by RTEEvents is treated in Section 4.2.2.3.

[rte_sws_2203] The RTE shall trigger the execution of Runnable Entities in accordance with the connected RTEEvent (see [RTE00072]).

[rte_sws_2204] The RTE-Generator shall reject configurations where not all RTE Events which can start a Runnable Entity are mapped to OS tasks (see [RTE00049] and [RTE00018]). The only exceptions are OperationInvokedEvents in case they are implemented by a direct function call.

[rte_sws_2207] The RTE shall respect the configured execution order of Runnable Entities within one OS task (see [RTE00070]).

With the information from Internal Behavior a part of the setup of the SW-Component within the RTE and the OS can already be configured. Furthermore, the information (description) of the structure (ports, interfaces) and the internal behavior of an AUTOSAR software component are sufficient for the *RTE Contract Phase*.

However, some detailed information is still missing and this is part of the Implementation description.

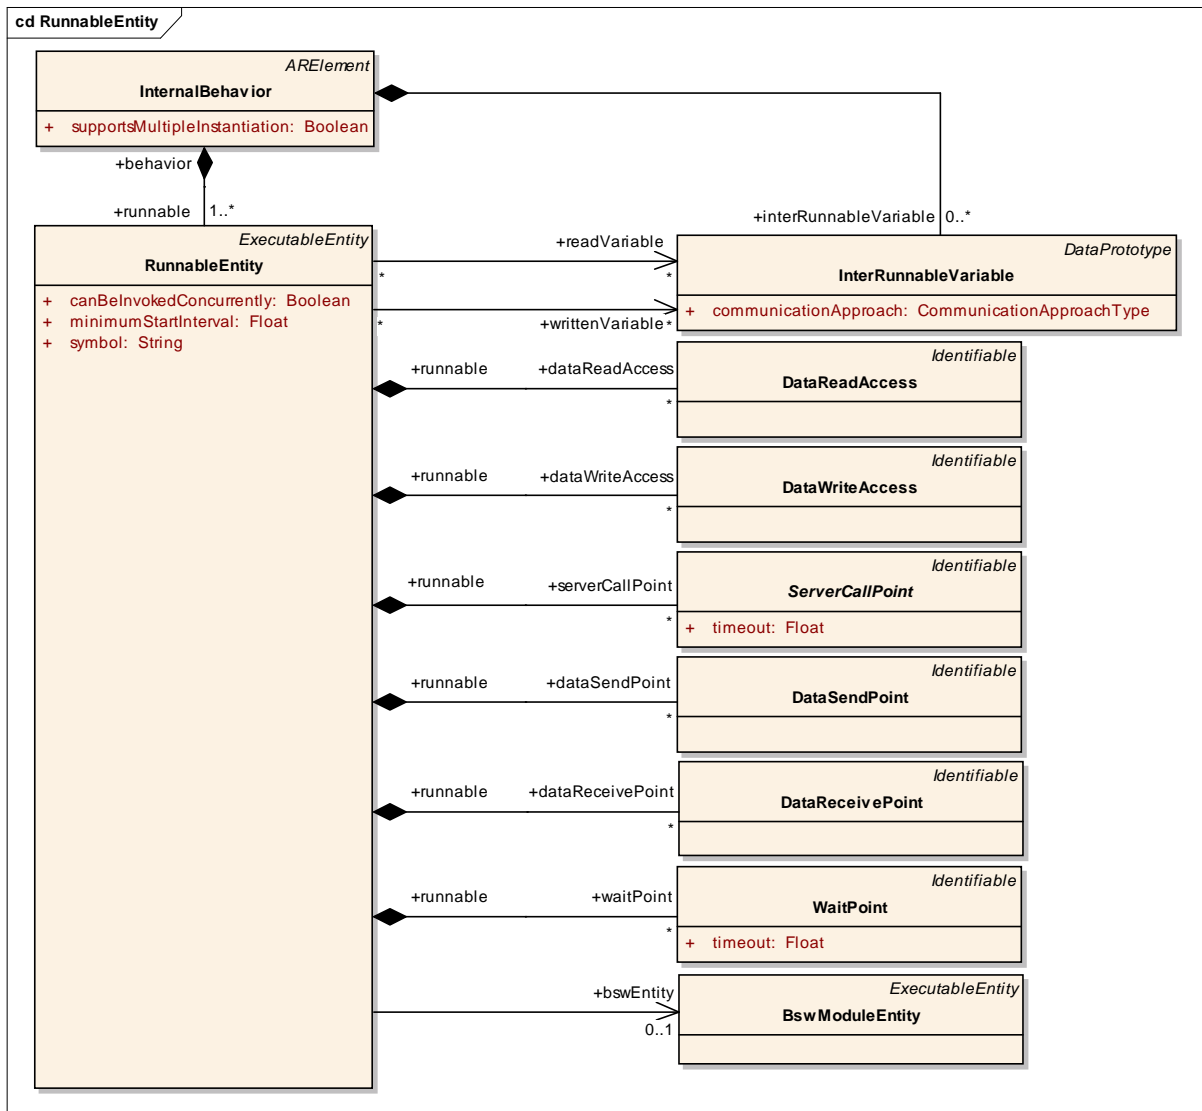


Figure 4.5: SW-Component runnable entity

4.1.2.4 Implementation

In the Implementation description an actual implementation of a SW-Component is described including the memory consumption (see Figure 4.6).

Note that the information from the Implementation part are only required for the *RTE Generation Phase*, if at all.

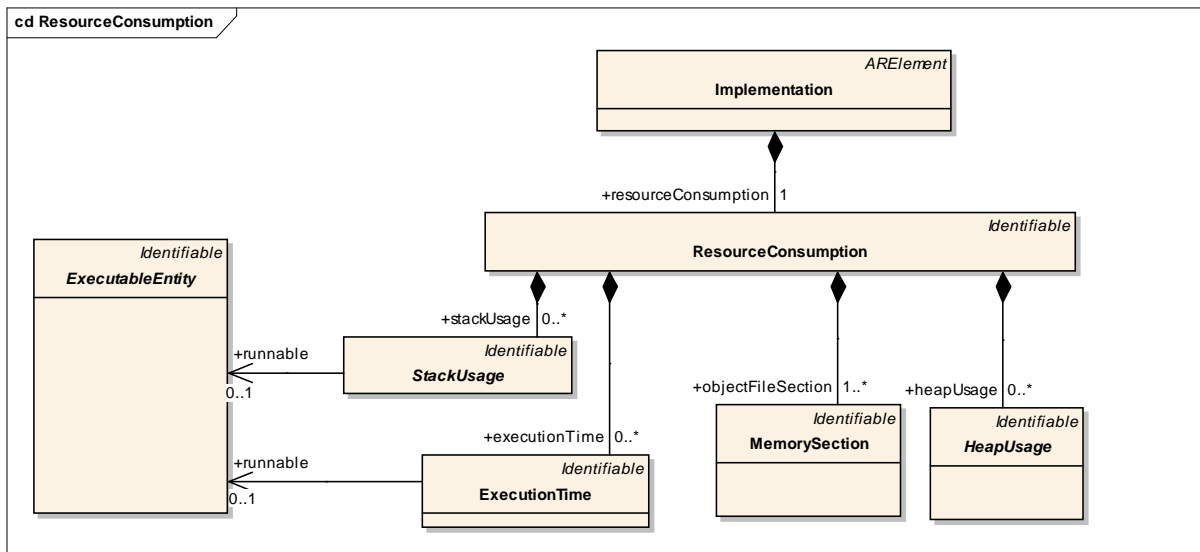


Figure 4.6: SW-Component resource consumption

4.1.3 Instantiation

4.1.3.1 Scope and background

Generally spoken, the term *instantiation* refers to the process of deriving specific instances from a model or template. But, this process can be accomplished on different levels of abstraction. Therefore, the instance of the one level can be the model for the next.

With respect to AUTOSAR four modeling levels are distinguished. They are referred to as the levels *M3* to *M0*.

The level *M3* describes the concepts used to derive an AUTOSAR meta model of level *M2*. This meta model at level *M2* defines a language in order to be able to describe specific attributes of a model at level *M1*, e.g., to be able to describe an specific type of an AUTOSAR software component. E.g., one part of the AUTOSAR meta model is called *Software Component Template* or *SW-C-T* for short and specified in [17]. It is discussed more detailed in section 4.1.2.

At level *M1* engineers will use the defined language in order to design components or interfaces or compositions, say to describe an specific *type* of a *LightManager*. Hereby, e.g., the descriptions of the (atomic) software components will also contain an internal behavior as well as an implementation part as mentioned in section 4.1.2.

Those descriptions are input for the RTE-Generator in the so-called 'Contract Phase' (see section 3.1). Out of this information specific APIs (in a programming language) to access ports and interfaces will be generated.

Software components generally consist of a set of Runnable Entities. They can now specifically be described in a programming language which can be referred to as "implementation". As one can see in section 4.1.2 these "implementation" then correspond

exactly to one implementation description as well as to one internal behavior description. However, they are still blueprints on *M1*.

M0 refers to a specific running instance on a specific car.

Objects derived from those specified component types can only be executed in a specific run time environment (on a specific target). The objects embody the real and running implementation and shall therefore be referred to as software component instances (on modeling level *M0*). E.g., there could be two component instances derived from the same component type *LightManager* on a specific *light controller* ECU each responsible for different lights. Making instances would mean here in first place, that it should be possible to distinguish them even though the objects are descended from the same model.

With respect to this more narrative description the *RTE* as the *run time environment* shall enable the process of instantiation. Thereby the term *instantiation* throughout the document shall refer to the process of deriving *M0* from *M1*. Therefore, this section will address the problems which can arise out of the instantiation process and will specify the needs for AUTOSAR components and the AUTOSAR RTE respectively.

[rte_sws_2000] The RTE-Generator shall be able to instantiate AUTOSAR software components out of an AUTOSAR software component description.

4.1.3.2 Concepts of instantiation

Regardless of the fact that the (aforementioned) instantiation of AUTOSAR software components can be generally achieved on a per-system basis, the RTE-Generator restricts its view to a per-ECU customization (see *rte_sws_5000*).

Generally, there are two different kinds of instantiations possible:

- single instantiation – which refers to the case where only *one* object or AUTOSAR software component instance will be derived out of the AUTOSAR software component description
- multiple instantiation – which refers to the case where *multiple* objects or AUTOSAR software component instances will be derived out of the AUTOSAR software component description

[rte_sws_2001] The RTE shall be able to instantiate one or more AUTOSAR software component instances out of a single AUTOSAR software component description.

[rte_sws_2008] The RTE-Generator shall evaluate the attribute *supportsMultipleInstantiation* of the *InternalBehavior* of an AUTOSAR software component description.

[rte_sws_2009] The RTE-Generator shall reject configurations where multiple instantiation is required, but the value of the attribute *supportsMultipleInstantiation* of the *InternalBehavior* of an AUTOSAR software component description is set to *FALSE*.

4.1.3.3 Single instantiation

Single instantiation refers to the easiest case of instantiation.

To be instantiated merely means that the code and the corresponding data of a particular RunnableEntity are embedded in a runtime context. In general, this is achieved by the context of an OS task (see example 4.1).

Example 4.1

Runnable entity R1 called out of a task context:

```

1      TASK (Task1) {
2          ...
3          R1 ();
4          ...
5      }
```

Since the single instance of the software component is unambiguous per se no additional concepts have to be added.

4.1.3.4 Multiple instantiation

[rte_sws_2002] Multiple objects instantiated from a single AUTOSAR software component (type) shall be identifiable without ambiguity.

There are two *principle* ways to achieve this goal –

- by code duplication (of runnable entities)
- by code sharing (of reentrant runnable entities)

For now it was decided to solely concentrate on code sharing and not to support code duplication.

[rte_sws_2017] Multiple instantiation shall be achieved by sharing code.

Multiple instances can share the same code, if the code is reentrant.

4.1.3.4.1 Reentrant code

In general, side effects can appear if the same code entity is invoked by different threads of execution running, namely tasks. This holds particularly true, if the invoked code entity inherits a state or memory by the means of static variables which are visible to all instances. That would mean that all instances are coupled by those static variables.

Thus, they affect each other. This would lead to data consistency problems on one hand. On the other – and that is even more important – it would introduce a new

communication mechanism to AUTOSAR and this is forbidden. AUTOSAR software components can only communicate via ports.

To be complete, it shall be noted that a calling code entity also inherits the reentrancy problems of its callee. This holds especially true in case of recursive calls.

4.1.3.4.2 Unambiguous object identification

[rte_sws_2015] The instantiated AUTOSAR software component objects shall be unambiguously identifiable by an *instance handle*, if multiple instantiation by sharing code is required.

4.1.3.4.3 Multiple instantiation and Per-instance memory

An AUTOSAR SW-C can define internal memory only accessible by a SW-C instance itself. This concept is called *PerInstanceMemory*. The memory can only be accessed by the runnable entities of this particular instance. That means in turn, other instances don't have the possibility to access this memory.

PerInstanceMemory API principles are explained in Section 5.2.5.

The API for *PerInstanceMemory* is specified in Section 5.6.12.

4.1.4 RTE and AUTOSAR Services

According to the AUTOSAR glossary [1] “an AUTOSAR service is a logical entity of the Basic Software offering general functionality to be used by various AUTOSAR software components. The functionality is accessed via standardized AUTOSAR interfaces”.

Therefore, AUTOSAR services provide standardized AUTOSAR Interfaces: ports typed by standardized *PortInterfaces*.

[rte_sws_2100] The RTE shall support the connection of AUTOSAR services only to AUTOSAR software-components located on the same ECU.

The RTE supports neither connections to AUTOSAR services located on remote ECUs nor connections between AUTOSAR services *rte_sws_2100*.

When connecting AUTOSAR service ports to ports of AUTOSAR software components the RTE maps standard RTE API calls to the symbols defined in the RTE input (i.e. XML) for the AUTOSAR service runnables of the BSW. The key technique to distinguish ECU dependent identifiers for the AUTOSAR services is called “port-defined argument values”, which is described in Section 4.3.2.4. Currently “port-defined argument values” are only supported for client-server communication. It is not possible to use a pre-defined symbol for sending or receiving data.

The RTE does not pass an instance handle to the C-based API of AUTOSAR services since the latter are single-instantiatable (see `rte_sws_3806`).

4.1.5 RTE and ECU Abstraction

The *ECU Abstraction* provides an interface to physical values for AUTOSAR software components. It abstracts the physical origin of signals (their pathes to the ECU hardware ports) and normalizes the signals with respect to their physical appearance (like specific values of current or voltage).

See the AUTOSAR ECU architecture in figure 4.2. From an architectural point of view the ECU Abstraction is part of the *Basic Software* layer and offers AUTOSAR interfaces to AUTOSAR software components. The *ECU Abstraction* is classified as firmware and will mostly interact with sensor and actuator software components.

Seen from the perspective of an RTE, regular AUTOSAR ports are connected. Without any restrictions all communication paradigms specified by the AUTOSAR Virtual Functional Bus (VFB) shall be applicable to the ports, interfaces and connections – sender-receiver just as well as client-server mechanisms.

However, ports of the ECU Abstraction shall always only be connected to ports of specific AUTOSAR software components: sensor or actuator software components. In this sense they are tightly coupled to a particular ECU Abstraction.

Furthermore, it must not be possible (by an RTE) to connect AUTOSAR ports of the ECU Abstraction to AUTOSAR ports of any AUTOSAR component located on a remote ECU (see `rte_sws_2051` and [RTE00136]).

This means, e.g., that sensor-related signals coming from the ECU Abstraction are always received by an AUTOSAR sensor component located on the same ECU. The AUTOSAR sensor component will then process the received signal and deploy it to other AUTOSAR components regardless of whether they are located on the same or any remote ECU. This applies to actuator-related signals accordingly, however, the opposite way around.

[rte_sws_ext_2054] The RTE-Generator expects only one instance of the ECU Abstraction.

[rte_sws_2050] The RTE-Generator shall generate a communication path between connected ports of AUTOSAR sensor or actuator software components and the ECU Abstraction in the exact same manner like for connected ports of AUTOSAR software components.

[rte_sws_2051] The RTE-Generator shall reject configurations which require a communication path from a AUTOSAR software component to an ECU Abstraction located on a remote ECU.

Further information about the ECU Abstraction can be found in the corresponding specification document [10].

4.1.6 RTE and Complex Device Driver

A Complex Device Driver has an AUTOSAR Interface, therefore the RTE can deal with the communication on the Complex Device Drivers ports. The Complex Device Driver is allowed to have code entities that are not under control of the RTE but yet still may use the RTE API (e.g. ISR2, BSW main functions).

4.2 RTE Implementation Aspects

4.2.1 Scope

This section describes some specific implementation aspects of an AUTOSAR RTE. It will mainly address

- the mapping of logical concepts (e.g., Runnable Entities) to technical architectures (namely, the AUTOSAR OS)
- the decoupling of pending interrupts (in the Basic Software) and the notification of AUTOSAR software components
- data consistency problems to be solved by the RTE

Therefore this section will also refer to aspects of the interaction of the AUTOSAR RTE and the two modules of the AUTOSAR Basic Software with standardized interfaces (see Figure 4.7):

- the module *AUTOSAR Operating System* [5, 11]
- the module *AUTOSAR COM* [3, 6]

Having a standardized interface means *first* that the modules do not provide or request services for/of the *AUTOSAR software components* located above the RTE. They do not have ports and therefore cannot be connected to the aforementioned AUTOSAR software components. AUTOSAR OS as well as AUTOSAR COM are simply invisible for them.

Secondly AUTOSAR OS and AUTOSAR COM are *used* by the RTE in order to achieve the functionality requested by the AUTOSAR software components. The AUTOSAR COM module is *used* by the RTE to route a signal over ECU boundaries, but this mechanism is hidden to the sending as well as to the receiving AUTOSAR software component. The AUTOSAR OS module is *used* by the RTE in order to properly schedule the single *Runnables* in the sense that the RTE-Generator generates *Task*-bodies which contain then the calls to appropriate Runnables.

In this sense the RTE shall also *use* the available means to convert interrupts to notifications in a task context or to guarantee data consistency.

With respect to this view the RTE is *thirdly not* the abstraction layer for AUTOSAR OS and AUTOSAR COM! Only the RTE offers the same *interface* to the AUTOSAR Soft-

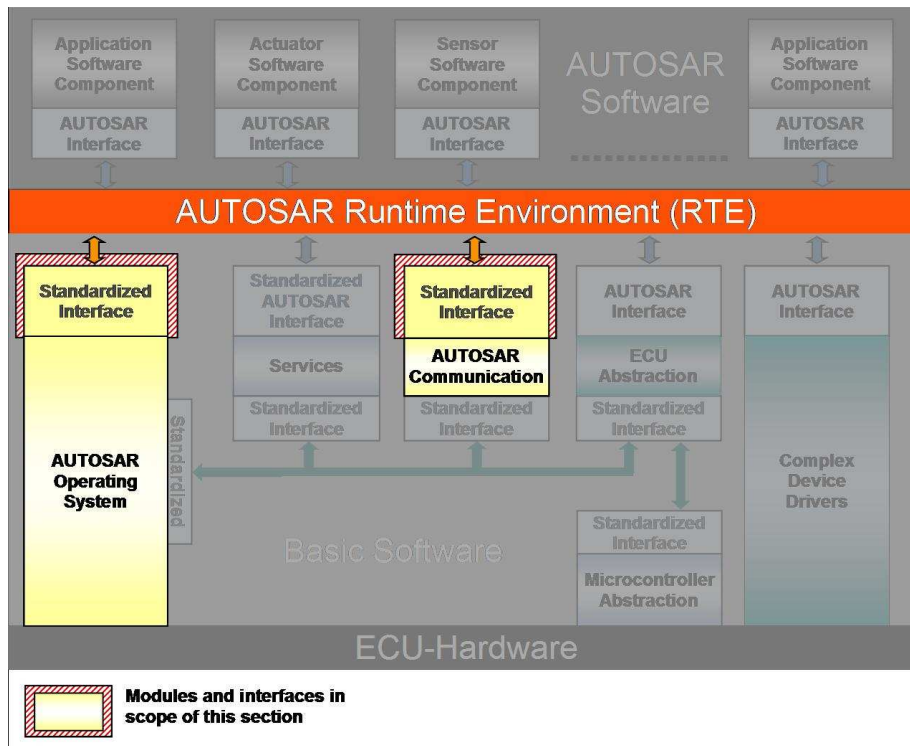


Figure 4.7: Scope of the section on Basic Software modules

ware Components like the VFB. For a specific ECU the RTE implements in conjunction with the modules of the Basis Software the entire functionality of the VFB (for that specific ECU). Hence, AUTOSAR OS and AUTOSAR COM are specific modules of a specific implementation of the VFB for a specific ECU. They shall be able to support the implementation of the VFB functionality, but the functionality of the modules are neither known by the AUTOSAR software components nor offered to them per se.

[rte_sws_2250] The RTE shall only use the AUTOSAR OS and AUTOSAR COM in order to provide the RTE functionality to the AUTOSAR components (see [RTE00020]).

[rte_sws_2251] The RTE-Generator shall construct task bodies for those tasks which contain Runnable Entities (see [RTE00049]).

The information for the construction of task bodies has to be given by the ECU Configuration description. The mapping of Runnable Entities to tasks is given as an input by the ECU Configuration description. The RTE-Generator does not decide on the mapping of Runnable Entities to tasks.

[rte_sws_2254] Missing input information for the RTE-Generator regarding the mapping of Runnable Entities to tasks or the construction of tasks bodies shall be taken as an invalid configuration and shall be rejected (see [RTE00049] and [RTE00018]).

4.2.2 OS

This chapter describes the interaction between the RTE and the AUTOSAR OS. The interaction is realized via the standardized interface of the OS - the AUTOSAR OS API. See Figure 4.7.

The OS is statically configured by the ECU-Configuration and not by the RTE generator. The RTE generator is not allowed to create tasks and other OS objects, which are necessary for the runtime environment. Also the mapping of runnable entities to tasks is not the job of the RTE generator. This mapping has to be done in a configuration step before, in the RTE-Configuration phase. The RTE generator is responsible for the generation of task bodies, which contain the calls for the runnable entities. The runnable entities themselves are OS independent and are not allowed to use OS service calls. The RTE has to encapsulate such calls via the standardized RTE API.

4.2.2.1 OS Objects

Tasks

- The RTE has to create the task bodies, which contain the calls of the runnable entities. Note that the term *task body* is used here to describe a piece of code, while the term *task* describes a configuration object of the OS.
- The RTE controls the task activation/resumption either directly by calling OS services like `SetEvent()` or `ActivateTask()` or indirectly by initializing OS alarms or starting Schedule-Tables for time-based activation of runnable entities. If the task terminates, the generated taskbody also contains the calls of `TerminateTask()` or `ChainTask()`.
- The RTE generator does **not** create tasks. The mapping of runnable entities to tasks is the input to the RTE generator and is therefore part of the RTE Configuration.
- The RTE configurator has to allocate the necessary tasks in the OS configuration.

OS applications

- The current AUTOSAR RTE SWS specification does not support memory protection.

Events

- The RTE may use OS Events for the implementation of the abstract RTEEvents.
- The RTE therefore may call the OS service functions `SetEvent()`, `WaitEvent()`, `GetEvent()` and `ClearEvent()`.
- The used OS Events are part of the input information of the RTE generator.

- The RTE configurator has to allocate the necessary events in the OS configuration.

Resources

- The RTE may use OS Resources (standard or internal) e.g. to implement data consistency mechanisms.
- The RTE may call the OS services `GetResource()` and `ReleaseResource()`.
- The used Resources are part of the input information of the RTE generator.
- The RTE configurator has to allocate the necessary resources (all types of resources) in the OS configuration.

Interrupt Processing

- An alternative mechanism to get consistent data access is disabling/enabling of interrupts. The AUTOSAR OS provides different service functions to handle interrupt enabling/disabling. The RTE may use these functions and must **not** use compiler/processor dependent functions for the same purpose.

Alarms

- The RTE may use Alarms for timeout monitoring of asynchronous client/server calls. The RTE is responsible for Timeout handling.
- The RTE may setup cyclic alarms for periodic triggering of runnable entities (runnable entity activation via `RTEEvent TimingEvent`)
- The used Alarms are part of the input information of the RTE generator.
- The RTE configurator has to allocate the necessary alarms in the OS configuration.

Schedule Tables

- The RTE may setup schedule tables for cyclic task activation (runnable entity activation via `RTEEvent TimingEvent`)
- The used schedule tables are part of the input information of the RTE generator.
- The RTE configurator has to allocate the necessary schedule tables in the OS configuration.

Memory Protection (SCC3/SCC4)

The current AutoSAR RTE specification does not support features of the AUTOSAR OS memory protection mechanisms. Nevertheless for future versions of the specification, the RTE is responsible to transfer the data of sender/receiver communication as well as for client-server communication over protection boundaries (OS applications).

Common OS features

Depending on the global scheduling strategy of the OS, the RTE can make decisions about the necessary data consistency mechanisms. E.g. in an ECU, where all tasks are non-preemptive - and as the result also the global scheduling strategy of the complete ECU is non-preemptive - the RTE may optimize the generated code regarding the mechanisms for data consistency.

Hook functions

The AUTOSAR OS Specification defines hook functions as follows:

A Hook function is implemented by the user and invoked by the operating system in the case of certain incidents. In order to react to these on system or application level, there are two kinds of hook functions.

- **application-specific:** Hook functions within the scope of an individual OS Application.
- **system-specific:** Hook functions within the scope of the complete ECU (in general provided by the integrator).

If no memory protection is used (scalability classes SCC1 and SCC2) only the system-specific hook functions are available.

The current AutoSAR RTE SWS specification does not support memory protection. Therefore, only the system-specific hooks are relevant. In the SRS the requirements to implement the system-specific hook functions are rejected [RTE00001], [RTE00101], [RTE00102] and [RTE00105]. The reason for the rejection is the system (ECU) global scope of those functions. The RTE is not the only user of those functions. Other BSW modules might have requirements to use hook functions as well. This is the reason why the RTE is not able to generate these functions without the necessary information of the BSW configuration.

It is intended that the implementation of the system specific hook functions is done by the system integrator and NOT by the RTE generator.

4.2.2.2 Runnable Entities

The following chapter describes the runnable entities, their categories and their task-mapping aspects. The prototypes of the functions implementing runnable entities are described in Chapter 5.7

Runnable entities are the schedulable parts of SW-Cs. With the exception of reentrant server runnables that are invoked via direct function calls, they have to be mapped to tasks. The mapping must be described in the ECU Configuration Description. This configuration - or just the RTE relevant parts of it - is the input of the RTE generator.

All runnable entities are activated by the RTE as a result an RTEEvent. Possible activation events are described in the meta-model by using RTEEvents (see Figure 4.8.

RTEEvents are described in the following chapter. If no RTEEvent is specified as StartOnEvent for the runnable entity, the runnable entity is never activated by the RTE.

The runnable entities are categorized as follows. **Category 1**

Category 1 runnable entities do not have *WaitPoints* and have to terminate in *finite* time. With respect to some constraints, category 1 runnable entities can be mapped to *Basic Tasks* of the AUTOSAR OS. The VFB Specification [13] distinguishes between Category 1A and Category 1B runnable entities. For mapping aspects, both sub categories can be handled equally and therefore the term *Category 1* is used instead.

Category 2

In contrast to category 1 runnable entities, runnable entities of category 2 always have at least one *WaitPoint* or they invoke a server and wait for the response to return (*SynchronousServerCallPoint*). Category 2 runnable entities usually have to be mapped to *Extended Tasks*, because only extended tasks provide the task state WAITING. The existence of at least one *WaitPoint* or of a *SynchronousServerCallPoint* classifies the runnable entity as a category 2 runnable.

Category 3

Runnable entities of category 3 are described in the VFB-Specification [13] in Chapter 4.5.4.4 but are currently out of scope of the RTE Specification. This restriction is also described in Section A.

4.2.2.3 RTE Events

The meta model describes the following RTE events.

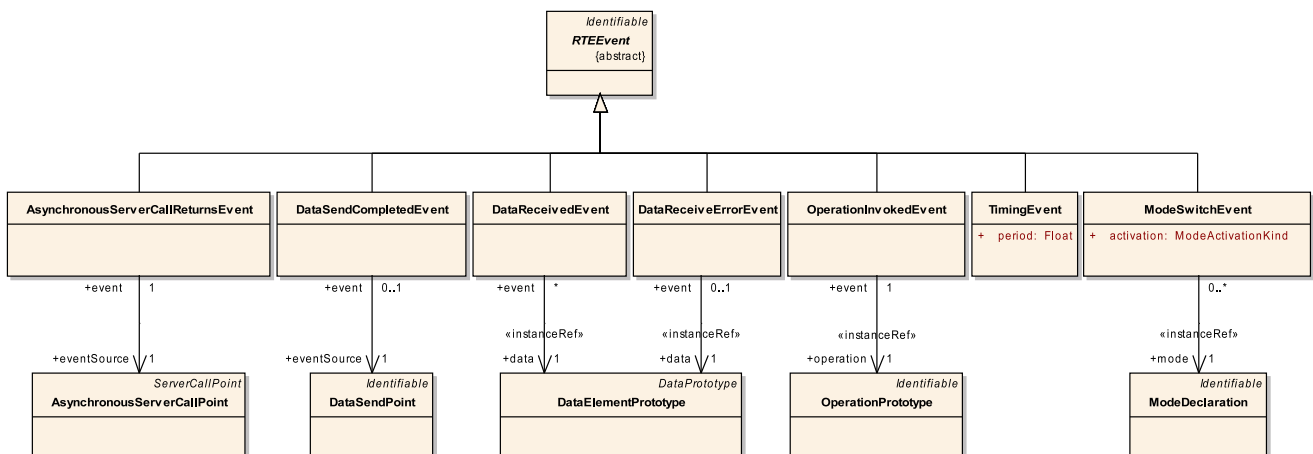


Figure 4.8: Different kinds of RTE-Events

T	TimingEvent
DR	DataReceivedEvent (S/R Communication only)
DRE	DataReceiveErrorEvent (S/R Communication only)
DSC	DataSendCompletedEvent (S/R Communication only)
OI	OperationInvokedEvent (C/S Communication only)
ASCR	AsynchronousServerCallReturnsEvent (C/S communication only)
MS	ModeSwitchEvent

According to the meta model it is possible that all kinds of RTEEvents can either

- 1.) **activate a runnable entity** or
- 2.) **wakeup a runnable entity at its waitpoints**

The meta model makes no restrictions. As a consequence RTE API functions would be necessary to set up the waitpoints for all kinds of RTEEvents.

Nevertheless in some cases it seems to make no sense to implement all possible combinations of the general meta model. E.g. setting up a waitpoint, which should be resolved by a cyclic TimingEvent. Therefore the RTE SWS of AUTOSAR Release 2.0 makes some restrictions, which are also described in Section A.

The meta model also allows, that the same runnable entity can be triggered by several RTEEvents. For the current approach of the RTE and restrictions see Section 4.2.5.

	T	DR	DRE	DSC	OI	ASCR	MS
Activation of runnable entity	x	x	x	x	x	x	x
Wakeup of waitpoint		x		x		x	

The table shows, that *activation of runnable entity* is possible for all kinds of RTEEvents. For runnable entity activation, no explicit RTE API is necessary. The RTE itself is responsible for the activation of the runnable entity depending on the configuration in the SW-C Description.

If the runnable entity contains a waitpoint, it can be resolved by the assigned RTE-Event(s). Entering the waitpoint requires an explicit call of a RTE API function. The RTE (together with the OS) has to implement the *Waitpoint* inside this RTE API.

The following list shows which RTE API function has to be called to set up waitpoints.

- DataReceivedEvent: `Rte_Receive()`
- DataSendCompletedEvent: `Rte_Feedback()`
- AsynchronousServerCallReturnsEvent: `Rte_Result()`

4.2.2.4 Mapping of runnable entities to tasks

One of the main requirements of the RTE is "Construction of task bodies" [RTE00049]. The necessary input information e.g. the mapping of runnable entities to tasks must be provided by the ECU configuration description.

The ECU configuration description (or an extract of it) is the input for the RTE-Generator (see Figure 3.3). It is also the purpose of this document to define the necessary input information. Therefore the following scenarios may help to derive requirements for the ECU-Configuration Template as well as for the RTE-generator itself.

Note: The scenarios do not cover all possible combinations.

The RTE-Configurator configures parts of the ECU-Configuration, e.g. the mapping of runnable entities to tasks. In this configuration process the RTE-Configurator also allocates those OS-objects (e.g. Tasks, Events, Alarms...) which are used in the generated RTE. The RTE-Configurator must be the **owner** of these configuration items. Other configurators, e.g. the OS Configurator, should not be able to change these settings.

Some figures for better understanding use the following conventions:




Figure 4.9: Element description

4.2.2.4.1 Scenario for mapping of runnable entities to tasks

The different properties of runnable entities with respect to data access and termination have to be taken into account when discussing possible scenarios of mapping runnable entities to tasks.

- Runnable entities using (implicit) `DataReadAccess/DataWriteAccess` have to terminate.
- Runnable entities using (implicit) `DataReadAccess/DataWriteAccess` are category 1 runnables (1A or 1B). Runnable entities of category 2 do not allow (implicit) `DataReadAccess/DataWriteAccess`.
- Runnable entities of category 1 can be mapped either to basic or extended tasks. (see next subsection).
- Runnable entities using at least one `Waitpoint` are of category 2.
- Runnables of category 2 that contain `WaitPoints` will be typically mapped to extended tasks.
- Runnables of category 2 that contain a `SynchronousServerCallPoint` generally have to be mapped to extended tasks.

- Runnables of category 2 that contain a SynchronousServerCallPoint can be mapped to basic tasks if no timeout monitoring is required.
- Runnables of category 2 that contain a SynchronousServerCallPoint can be mapped to basic tasks if the server runnable is invoked directly and is itself of category 1.

Note that the runnable to task mapping scenarios supported by a particular RTE implementation might be restricted.

4.2.2.4.1.1 Scenario 1

Runnable entity category 1A: "runnable1"

- Ports: only S/R with DataReadAccess / DataWriteAccess
- RTEEvents: TimingEvent
- no sequence of runnable entities specified
- no explicit DataSendPoint
- no WaitPoint

Possible mappings of "runnable1" to tasks:

Basic Task

If only one of those kinds of runnable entities is mapped to a task (task contains only one runnable entity), or if multiple runnable entities with the same cycletime are mapped to the same task, a basic task can be used. In this case, the execution order of the runnable entities within the task is necessary. In case the runnable entities have different cycletimes, the RTE has to provide the glue-code to guarantee the correct call cycle of each runnable entity.

The ECU-Configuration-Template has to provide the sequence of runnable entities mapped to the same task, see `rte_sws_in_0014`.

Figure 4.10 shows the possible mappings of runnable entities into a basic task. If and only if a sequence order is specified, more than one runnable entity can be mapped into a basic task.

Extended Task

If more than one runnable entity is mapped to the same task and the special condition (same cycletime) does not fit, an extended task is used.

If an extended task is used, the entry points to the different runnable entities might be distinguished by evaluation of different OS events. In the scenario above, the different cycletimes may be provided by different OS alarms. The corresponding OS events have to be handled inside the task body. Therefore the RTE-generator needs for each task the number of assigned OS Events and their names.




Figure 4.10: Mapping of Category 1 runnable entities to Basic Tasks

The ECU-Configuration has to provide the OS events assigned to the RTEEvents triggering the runnable entities that are mapped to an extended task, see `rte_sws_in_0039`.

Figure 4.11 shows the possible mapping of the multiple runnable entities of category 1 into an Extended Task. Note: The Task does not terminate.




Figure 4.11: Mapping of Category 1 runnable entities to Extended Tasks

For both, basic tasks and extended tasks, the ECU-Configuration must provide the name of the task.

The ECU-Configuration has to provide the name of the task, see `rte_sws_in_5012`.

The ECU-Configuration has to provide the task type (BASIC or EXTENDED), which can be determined from the presence or absence of OS Events associated with that task, see `rte_sws_in_0040`.

4.2.2.4.1.2 Scenario 2

Runnable entity category 1B: "runnable2"

- Ports: S/R with DataSendPoints.
- RTEEvents: TimingEvent
- no sequence of runnables specified
- no WaitPoint

Possible mappings of "runnable2" to tasks:

The following figure shows the different mappings:

- One category 1B runnable
- More than one category 1B runnable mapped to the same basic task with a specified sequence order
- More than one category 1B runnable mapped into an extended task

The gluecode to realize the DataReadAccess and DataWriteAccess respectively before entering the runnable and after exiting is not necessary.




Figure 4.12: Mapping of Category 1 runnable entities using no DataReadAccess / DataWriteAccess

4.2.2.4.1.3 Scenario 3

Runnable entity category 1A: "runnable3"

- Ports: S/R with DataReadAccess / DataWriteAccess
- RTEEvents: Runnable is activated by a DataReceivedEvent

- no sequence of runnables specified
- no DataSendPoint
- no WaitPoint

There is no difference between Scenario 1. Only the RTEEvent that activates the runnable entity is different.

4.2.2.4.1.4 Scenario 4

Runnable entity category 2: "runnable4"

- Ports: S/R with DataReceivePoint and WaitPoint (blocking read)
- RTEEvents: WaitPoint referencing a DataReceivedEvent
- no sequence of runnables specified

Runnable is activated by an arbitrary RTEEvent (e.g. by a TimingEvent). When the runnable entity has entered the WaitPoint and the DataReceivedEvent occurs, the runnable entity resumes execution.

The runnable has to be mapped to an extended task. Normally each category 2 runnable has to be mapped to its own task. Nevertheless it is not forbidden to map multiple category 2 runnable entities to the same task, though this might be restricted by an RTE generator. Mapping multiple category 2 runnable entities to the same task can lead to big delay times if e.g. a WaitPoint is resolved by the incoming RTEEvent, but the task is still waiting at a different WaitPoint.




Figure 4.13: Mapping of Category 2 runnable entities to Extended Tasks

4.2.2.4.1.5 Scenario 5

There are two runnable entities implementing a client (category 2) and a server for synchronous C/S communication and the timeout attribute of the ServerCallPoint is 0.

There are two ways to invoke a server synchronously:

- Simple function call for intra-ECU C/S communication if the canBeInvokedConcurrently attribute of the server runnable is set and if the server runnable is of

category 1. In that case the server runnable is executed in the same task context (same stack) as the client runnable that has invoked the server. The client runnable can be mapped to a basic task.

- The server runnable is mapped to its own task. If the `canBeInvokedConcurrently` attribute is not set, the server runnable must be mapped to a task.

If the implementation of the synchronous server invocation does not use OS events, the client runnable can be mapped to a basic task and the task of the server runnable must have higher priority than the task of the client runnable. Furthermore, the task to which the client runnable is mapped must be preemptable. This has to be checked by the RTE generator. Activation of the server runnable can be done by `ActivateTask()` for a basic task or by `SetEvent()` for an extended task. In both cases, the task to be activated must have higher priority than the task of the client runnable to enforce a task switch (necessary, because the server invocation is synchronous).

4.2.2.4.1.6 Scenario 6

There are two runnable entities implementing a client (category 2) and a server for synchronous C/S communication and the timeout attribute of the `ServerCallPoint` is greater than 0.

There are again two ways to invoke a server synchronously:

- Simple function call for intra-ECU C/S communication if the `canBeInvokedConcurrently` attribute of the server runnable is set and the server is of category 1. In that case the server runnable is executed in the same task context (same stack) as the client runnable that has invoked the server and no timeout monitoring is performed (see `rte_sws_3768`). In this case the client runnable can be mapped to a basic task.
- The server runnable is mapped to its own task. If the `canBeInvokedConcurrently` attribute is not set, the server runnable must be mapped to a task.

If the implementation of the timeout monitoring uses OS events, the task of the server runnable must have lower priority than the task of the client runnable and the client runnable must be mapped to an extended task. Furthermore, both tasks must be preemptable¹. This has to be checked by the RTE generator. The notification that a timeout occurred is then notified to the client runnable by using an OS Event. In order for the client runnable to immediately react to the timeout, a task switch to the client task must be possible when the timeout occurs.

¹Strictly speaking, this restriction is not necessary for the task to which the client runnable is mapped. If OS events are used to implement the timeout monitoring and the notification that the server is finished, the RTE API implementation generally uses the OS service `WaitEvent`, which is a point of rescheduling.

4.2.2.4.1.7 Scenario 7

Runnable entity category 2: "runnable7"

- Ports: only C/S with AsynchronousServerCallPoint and WaitPoint
- RTEEvents: AsynchronousServerCallReturnsEvent (C/S communication only)
- no sequence of runnables specified

The mapping scenario for "runnable7", the client runnable that collects the result of the asynchronous server invocation, is similar to Scenario 4.

4.2.2.5 Activation Offset for runnable

In order to allow optimizations (smooth cpu load, mapping of runnables with different periods in the same task to avoid data sharing, etc.), the RTE has to handle the activation offset information from a task shared reference point only for time trigger runnables. The maximum period of a task can be calculated automatically as the greatest common divisor (GCD) of all runnables period and offset. It is assumed that the runnables worst case execution is less than the GCD. In case of the worst case execution is greater than the GCD, the behavior becomes undefined.

[rte_sws_7000] The RTE shall respect the configured activation offset of runnable entities mapped within one OS task.

[rte_sws_ext_7001] The runnables worst case execution time shall be less than the GCD of all runnables period and offset in activation offset context for runnables.

Example 1:

This example describes 3 runnables mapped in one task with an activation offset defined for each runnables.

Runnable	Period	Activation Offset
R1	100ms	20ms
R2	100ms	60ms
R3	100ms	100ms

Table 4.1: Runnables timings

The runnables R1, R2 and R3 are mapped in the task T1 at 20 ms which is the GCD of all runnables period and activation offset.




Figure 4.14: Example of activation offset for runnables

Example 2:

This example describes 4 runnables mapped in one task with an activation offset and position in task defined for each runnables.

Runnable	Period	Position in task	Activation Offset
R1	50ms	1	0ms
R2	100ms	2	0ms
R3	100ms	3	70ms
R4	50ms	4	20ms

Table 4.2: Runnables timings with position in task

The runnables R1, R2,R3 and R4 are mapped in the task T1 at 10 ms which is the GCD of all runnables period and activation offset.




Figure 4.15: Example of activation offset for runnables with position in task

4.2.2.6 Activation and Start of Runnable Entities

This section defines the activation of a runnable entity by using a state machine.

The main principles for the activation of runnables are:

- runnables are activated by RTE events
- runnable activations are not queued (with exception of OperationInvokedEvent).
If a runnable is activated due to several DataReceivedEvents of DataElements with isQueued=true, it is the responsibility of the runnable to dequeue all queued data.
- A 'minimum start interval' will delay the activation of a runnable to prevent that a runnable is started more than once within the 'minimum start interval'.

Each runnable has its own state machine to describe all necessary states and transitions between a suspended and a running runnable. The runnable state machine is shown in Fig. 4.16.

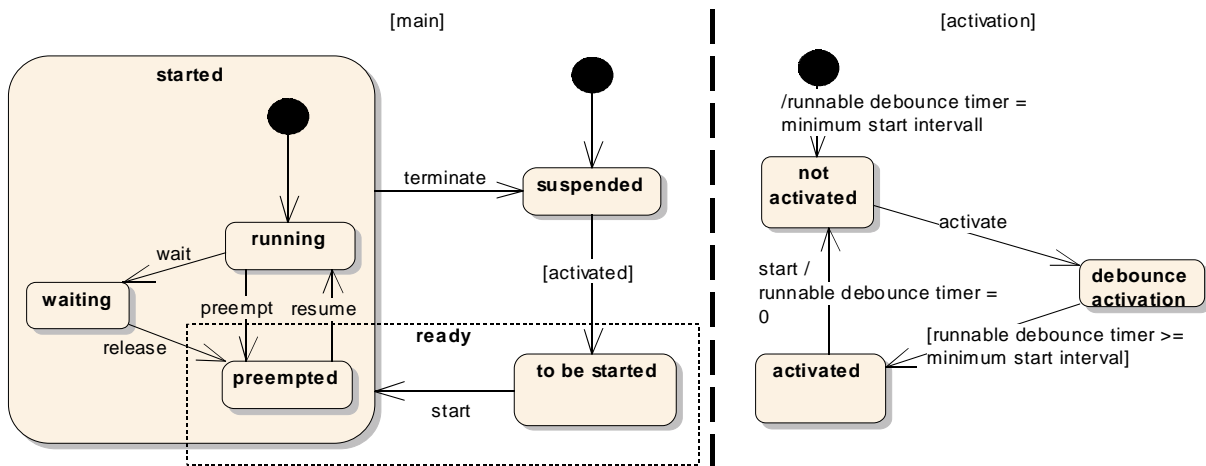


Figure 4.16: State machine of a runnable entity (not a server runnable) Note: the runnable debounce timer is an increasing timer. It is local to the runnable. The runnable debounce timer is just a concept for the specification, not for the implementation.

The state machine of a runnable is not identical to that of the task containing the runnable, but there are dependencies between them. E.g., the runnable can only be 'running' when the corresponding task is 'running'.

Table 4.3 describes all runnable states in detail. The runnable state machine is split in two sub state machines. The runnable main states describe the real state of the runnable, while the activation states describe the state of the pending activations by RTE events.

runnable state	description
runnable main states	
suspended	The runnable is not started and there is no pending request to start the runnable.

to be started	The runnable is activated but not yet started. Entering the to be started state, usually implies the activation of a task that starts the runnable. The runnable stays in the 'to be started' state, when the task is already running until the gluecode of the task actually starts the execution of the runnable.
running	The runnable code is being executed. A task that contains the runnable is running.
waiting	A task containing the runnable is waiting at a WaitPoint within the runnable.
preempted	A task containing the runnable is preempted from executing the runnable code.
started	'started' is the super state of 'running', 'waiting' and 'preempted' between start and termination of the runnable.
runnable activation states	
not activated	No RTE event requires the activation of the runnable.
debounce activation	One or more RteEvents with a startOnEvent relation to the runnable have occurred, but the debounce timer has not yet exceeded the minimum start interval. The activation will automatically advance to activated, when the debounce timer reaches the minimum start interval.
activated	One or more RteEvents with a startOnEvent to the runnable have occurred, and the debounce timer has exceeded the minimum start interval. While the activated state is active, the main state of the runnable automatically advances from the suspended to the 'to be started' state. When the runnable starts, the activation will be reset to 'not activated'.

Table 4.3: States defined for each runnable.

Note: For tasks, the equivalent state machine does not distinguish between preempted and to be started. They are subsumed as 'ready'

[rte_sws_2697] The activation of runnable entities (except for `server` runnables) shall behave as described by the runnable state machine in Fig. 4.16 and Table 4.3.

The following examples in Fig. 4.17 and Fig. 4.18 show the different timing situations of the runnables with or without a minimum start interval. The minimum start interval can reduce the number of activations by collecting more activating RTE Events within that interval. No activation will be lost, the activations are just delayed to keep the minimum start interval.

When a data received event activates a runnable when it is still running, it might be that the data is already dequeued during the current execution of the runnable. Still, the runnable will be started again. So, it is possible that a runnable that is activated by a data received event finds an empty receive queue.

A `server` runnable is exclusively activated by `OperationInvokedEvents` and implements the `server` in client server communication. In some cases, the client server




Figure 4.17: This figure illustrates the activation of a runnable without minimum start interval. The started state of the runnable main states and the activated state of the runnable activation states are shown. Each flash indicates the occurrence of an RTE event. The runnable can only be activated once. The activation is not queued. The runnable can be activated again when it is still started.




Figure 4.18: This figure illustrates the activation of a runnable with a minimum start interval. The started state of the runnable main states and the debounce activation and activated states of the runnable activation states are shown. Each flash indicates the occurrence of an RTE event. The red arrows indicate the minimum start interval after each start of the runnable. An RTE event within this minimum start interval leads to the debounce activation state. When the minimum start interval ends, the debounce activation state changes to the activated state.

communication is implemented by RTE as a direct function call of the server by the client. In this case, a minimum start distance is not applicable.

[rte_sws_2698] RTE shall not use the minimum start interval for `server` runnables.

When a server runnable is called by direct function call, the states 'to be started', 'debounce activation', and 'activated' are passed immediately. If a server runnable can be invoked concurrently and is invoked by different clients using direct function call, each client's concurrent execution of the server runnable has its own state machine.

[rte_sws_2699] When RTE implements server serialization, the activations of the server shall be queued according to Fig. 4.19.b.

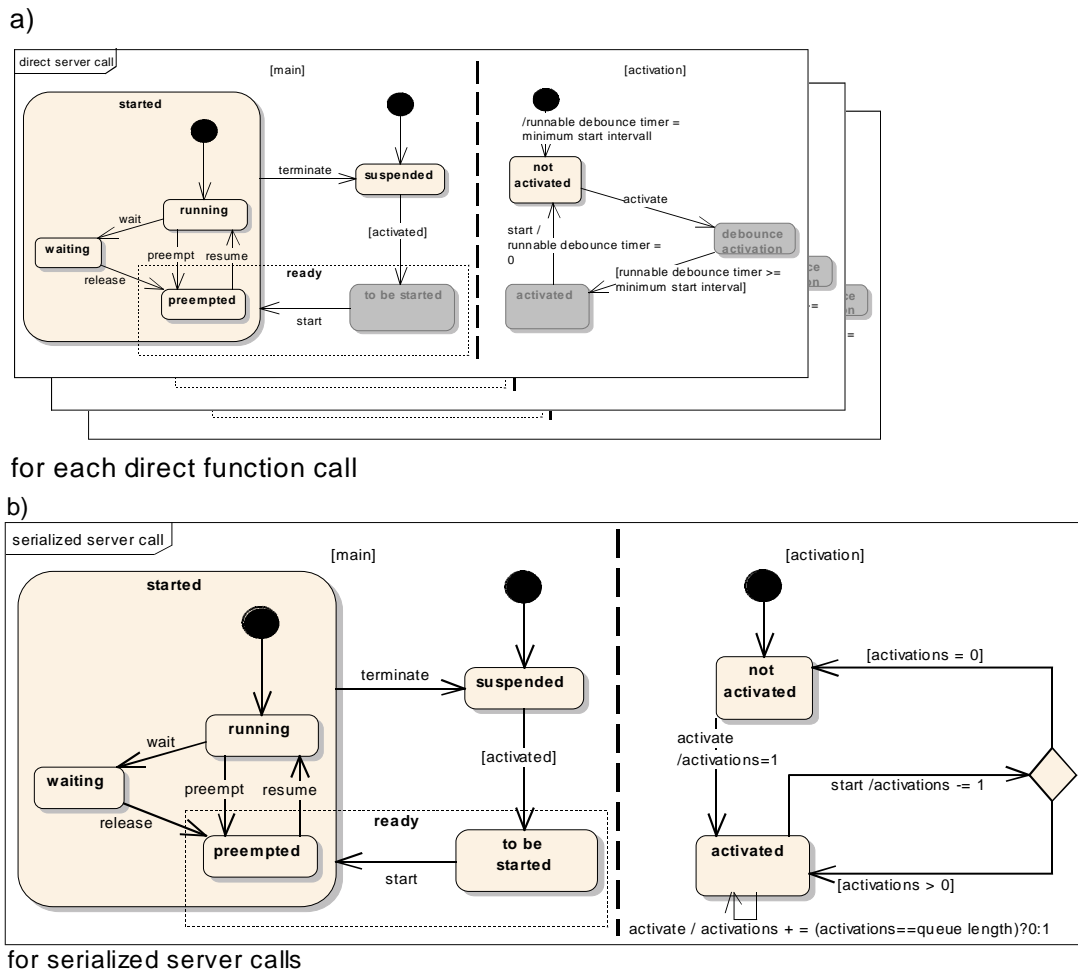


Figure 4.19: State machines of a server runnable. (a) For each client where RTE implements the server call as a direct function call, the server runnable has a separate state machine as displayed in the upper state charts. States that are directly passed are grayed out. (b) The client server transactions that are queued by RTE require a queuing of the activations as displayed in the lower state chart

4.2.3 Interrupt decoupling and notifications

4.2.3.1 Basic notification principles

Several BSW modules exist which contain functionality which is not directly activated, triggered or called by AUTOSAR software-components but by other circumstances, like digital input port level changes, complex driver actions, CAN signal reception, etc. In most cases interrupts are a result of those circumstances. For a definition of interrupts, see the VFB [13].

Several of these BSW functionalities create situations, signalled by an interrupt, when AUTOSAR SW-Cs have to be involved. To inform AUTOSAR software components of those situations, runnables in AUTOSAR software components are activated by notifications. So interrupts that occur in the basic software have to be transformed into notifications of the AUTOSAR software components. Such a transformation has to take

place at RTE level **at the latest!** Which interrupt is connected to which notification is decided either during system configuration/generation time or as part of the design of Complex Device Drivers or the Microcontroller Abstraction Layer.

This means that runnables in AUTOSAR SW-Cs have to be activated or "waiting" cat2 runnables in extended tasks have to be set to "ready to run" again. In addition some event specific data may have to be passed.

There are two different mechanisms to implement these notifications, depending on the kind of BSW interfaces.

1. **BSW with Standardized interface.** Used with COM and OS.
Basic-SW modules with Standardized interfaces cannot create RTEEvents. So another mechanism must be chosen: "**callbacks**"
The typical callback realization in a C/C++ environment is a function call.
2. **BSW with AUTOSAR interface:** Used in all the other BSW modules.
Basic-SW modules with AUTOSAR-Interfaces have their interface specified in an AUTOSAR BSW description XML file which contains signal specifications according to the AUTOSAR specification. The BSW modules can employ RTE API calls like `Rte.Send` – see 5.6.4). **RTEEvents** may be connected with the RTE API calls, so realizing AUTOSAR SW-C activation.

Note that an AUTOSAR software component can send a notification to another AUTOSAR software component or a BSW module only via an AUTOSAR interface.

4.2.3.2 Interrupts

The AUTOSAR concept as stated in the VFB specification [13] does not allow AUTOSAR software components to run in interrupt context. Only the Microcontroller Abstraction Layer, Complex Device Drivers and the OS are allowed to directly interact with interrupts and implement interrupt service routines (see Requirement BSW164). This ensures hardware independency and determinism.

If AUTOSAR software components were allowed to run in interrupt context, one AUTOSAR software component could block the entire system schedule for an unacceptably long period of time. But the main reason is that AUTOSAR software components are supposed to be independent of the underlying hardware so that exchangeability between ECUs can be ensured. The schedule of an ECU is more predictable and better testable if the timing effects of interrupts are restricted to the basic software of that ECU.

Furthermore, AUTOSAR software components are not allowed to explicitly block interrupts as a means to ensure data consistency. They have to use RTE functions for this purpose instead, see Section 4.2.4.

4.2.3.3 Decoupling interrupts on RTE level

Runnables in AUTOSAR SW-Cs may be running as a consequence of an interrupt but **not** in interrupt context, which means not within an interrupt service routine! Between the interrupt service routine and an AUTOSAR SW-C activation there must always be a decoupling instance. AUTOSAR SW-C runnables are only executed in the context of tasks.

The decoupling instance is latest the RTE. For the RTE there are several options to realize the decoupling of interrupts. Which option is the best depends on the configuration and implementation of the RTE, so only examples are given here.

Example 1:

Situation:

- An interrupt routine calls an RTE callback function

Intention:

- Start a runnable

RTE job:

- RTE starts a task containing the runnable activation code by using the "Activate-Task()" OS service call.
- Other more sophisticated solutions are possible, e.g. if the task containing the runnable is activated periodically.

Example 2:

Situation:

- An interrupt routine calls an RTE callback function

Intention:

- Make a runnable wake up from a wait point

RTE job:

- RTE sets an OS event

These scenarios described in the examples above not only hold for RTE callback functions but for other RTE API functions as well.

[rte_sws_3600] The RTE shall prevent runnable entities of AUTOSAR software-components to run in interrupt context.

4.2.3.4 RTE and interrupt categories

Since category 1 interrupts are not under OS control the RTE has absolutely no possibility to influence their execution behavior. So no category 1 interrupt is allowed to reach RTE. This is different for interrupt of category 2.

[rte_sws_3594] Only interrupt category 2 can use RTE services.

4.2.3.4.1 Interrupt decoupling for COM

COM callbacks are used to inform the RTE about something that happened independently of any RTE action. This is often interrupt driven, e.g. when a data item has been received from another ECU or when a S/R transmission is completed.

It is the RTE's job e.g. to create RTEEvents from the interrupt.

[rte_sws_3530] The RTE has to provide callback functions to allow COM to signal COM events to the RTE.

[rte_sws_3531] The RTE has to support runnable activation by COM callbacks.

[rte_sws_3532] The RTE has to support cat2 runnables to wake up from a wait point as a result of COM callbacks.

See RTE callback API in chapter 5.9.

4.2.4 Data Consistency

4.2.4.1 General

Concurrent accesses to shared data memory can cause data inconsistencies. In general this must be taken into account when several code entities accessing the same data memory are running in tasks with different priority levels - in other words when systems using parallel (or quasi parallel) execution of code are designed. More general: Whenever task context-switches occur and data is shared between tasks, data consistency is an issue.

AUTOSAR systems use operating systems according to the AUTOSAR-OS specification which is derived from the OSEK-OS specification. The Autosar OS specification defines a priority based scheduling to allow event driven systems. This means that tasks with higher priority levels are able to interrupt (preempt) tasks with lower priority level.

The "lost update" example in Figure 4.20 illustrates the problem for concurrent read-modify-write accesses:

There are two tasks. Task A has higher priority than task B. A increments the com-




Figure 4.20: Data inconsistency example - lost update

only accessed counter X by 2, B increments X by 1. So in both tasks there is a read (step1) – modify (step2) – write (step3) sequence. If there are no atomic accesses (fully completed read-modify-write accesses without interruption) the following can happen:

1. Assume X=5.
2. B makes read (step1) access to X and stores value 5 in an intermediate store (e.g. on stack or in a CPU register).
3. B cannot continue because it is preempted by A.
4. A does its read (step1) – modify (step2) – write (step3) sequence, which means that A reads the actual value of X, which is 5, increments it by 2 and writes the new value for X, which is 7. ($X=5+2$)
5. A is suspended again.
6. B continues where it has been preempted: with its modify (step2) and write (step3) job. This means that it takes the value 5 from its internal store, increments it by one to 6 and writes the value 6 to X ($X=5+1$).
7. B is suspended again.

The correct result after both Tasks A and B are completed should be X=8, but the update of X performed by task A has been lost.

4.2.4.2 Communications to look at

In AUTOSAR systems the RTE has to take care that a lot of the communication is not corrupted by data consistency problems. RTE Generator has to apply suitable means if required.

The following communication mechanisms can be distinguished:

- Intra ECU communication within one AUTOSAR SW-C:
Communication between Runnables of one AUTOSAR SW-C running in different task contexts where communication between these Runnables takes place via commonly accessed data. If the need to support data consistency by the RTE exists it must be specified by using the concepts of "ExclusiveAreas" or "Inter-RunnableVariables" only.
- Intra-ECU communication between AUTOSAR SW-Cs:
Sender/Receiver (S/R) communication between Runnables of different AUTOSAR SW-Cs using *implicit* or *explicit* data exchange can be realized by the RTE through commonly accessed RAM memory areas. Data consistency in Client/Server (C/S) communication can be put down to the same concepts as S/R communication. Data access collisions must be avoided. The RTE is responsible for guaranteeing data consistency.
- Intra-ECU communication between AUTOSAR SW-Cs and BSW modules with AUTOSAR interfaces:
Principally the same as above: Sender/Receiver (S/R) communication between AUTOSAR SW-Cs and BSW modules using *implicit* or *explicit* data exchange can be realized by the RTE through shared RAM memory areas. Data consistency in Client/Server (C/S) communication can be put down to the same concepts as S/R communication. Data access collisions must be avoided. Again, the RTE has to guarantee data consistency!
- Inter ECU communication
COM has to guarantee data consistency for communication between ECUs on complete path between the COM modules of different ECUs. The RTE on each ECU has to guarantee that no data inconsistency might occur when it invokes COM send respectively receive calls supplying respectively receiving data items which are concurrently accessed by application via RTE API call, especially when queueing is used since the queues are provided by the RTE and not by COM.

[rte_sws_3514] The RTE has to guarantee data consistency for communication via AUTOSAR interfaces.

4.2.4.3 Concepts

In the AUTOSAR SW-C Template [17] chapter "Interaction between runnables within one component", the concepts of

1. ExclusiveAreas (*see section 4.2.4.5 below*)
2. InterRunnableVariables (*see section 4.2.4.6 below*)

are introduced to allow the user (SW-Designer) to specify where the RTE shall guarantee data consistency for AUTOSAR SW-C internal communication and execution circumstances. This is discussed in more detail in next sections.

The AUTOSAR SW-C template specification [17] also states that AUTOSAR SW-Cs may define **PerInstanceMemory**, allowing reservation of static (permanent) need of global RAM for the SW-C. Nothing is specified about the way Runnables might access this memory. RTE only provides a reference to this memory (*see section 5.6*) but doesn't guarantee data consistency for it.

The creator of an AUTOSAR SW-C has to take care by himself that accesses to RAM reserved as PerInstanceMemory out of Runnables running in different task contexts don't cause data inconsistencies. On the other hand this provides more freedom in using the memory.

4.2.4.4 Mechanisms to guarantee data consistency

ExclusiveAreas and InterRunnableVariables are only mentioned in association with AUTOSAR SW-C internal communication. Nevertheless the data consistency mechanisms behind can be applied to communication between AUTOSAR SW-Cs or between AUTOSAR SW-Cs and BSW modules too. Everywhere where the RTE has to guarantee data consistency.

The data consistency guaranteeing mechanisms listed here are derived from AUTOSAR SW-C Template and from further discussions. There might be more.

The RTE has the responsibility to apply such mechanisms if required. The details how to apply the mechanisms are left open to the RTE supplier.

Mechanisms:

- **Sequential scheduling strategy**

The activation code of Runnables is sequentially placed in one task so that no interference between them is possible because one Runnable is only activated after the termination of the other. Data consistency is guaranteed.

- **Interrupt blocking strategy**

Interrupt blocking can be an appropriate means if collision avoidance is required for a very short amount of time. This might be done by disabling respectively suspending all interrupts or - if hardware supports it - only of some interrupt levels. In general this mechanism must be applied with care because it might influence SW in tasks with higher priority too and the timing of the complete system.

- **Usage of OS resources**

Usage of OS resources. Advantage in comparison to Interrupt blocking strategy is that less SW parts with higher priority are blocked. Disadvantage is that implementation might consume more resources (code, runtime) due to the more sophisticated mechanism.

- **Task blocking strategy**

Mutual task preemption is prohibited. This might be reached e.g. by assigning same priorities to affected tasks, by assigning same internal OS resource to affected tasks or by configuring the tasks to be non-preemptive.

- **Cooperative Runnable placement strategy**

The principle is that tasks containing Runnables to be protected by "Cooperative Runnable placement strategy" are not allowed to preempt other tasks also containing Runnables to be protected by "Cooperative Runnable placement strategy" when one of the Runnables to protect is active - but are allowed between Runnable executions. The RTE's job is to create appropriate task bodies and use OS services or other mechanisms to achieve the required behavior.

To point out the difference to "Task blocking strategy":

In "Task blocking strategy" no task containing Runnables with access to the ExclusiveArea at all is allowed to preempt another task containing Runnables with access to same ExclusiveArea. In "Cooperative Runnable placement strategy" this task blocking mechanism is limited to tasks defined to be within same cooperative context.

Example to explain the cooperative mechanism:

- Runnables R2 and R3a are marked to be protected by cooperative mechanism.
- Runnables R1, R3b and R4 have no cooperative marking.
- R1 is activated in Task T1, R2 is activated in Task T2, R3a is activated in Task T3a, R3b is activated in Task T3b, R4 is activated in Task T4.
- Task priorities are: $T4 > T3a > T2 > T1$, T3b has same priority as T3a

This setup results in this behavior:

- T4 can always preempt all other tasks (Higher prio than all others).
- T3b can preempt T2 (higher prio of T3b, no cooperative restriction)
- T3a cannot preempt T2 (Higher prio of T3a but same cooperative context). So data access of Runnable R2 to common data cannot interfere with data access by Runnable R3a. Nevertheless if both tasks T3a and T2 are ready to run, it's guaranteed that T3a is running first.
- T1 can never preempt one of the other tasks because of lowest assigned prio.

- **Copy strategy**

Idea: The RTE creates copies of data items so that concurrent accesses in different task contexts cannot collide because some of the accesses are redirected to the copies.

How it can work:

- Application for **read conflicts**:

For all readers with lower priority than the writer a *read copy* is provided.

Example:

There exist Runnable R1, Runnable R2, data item X and a copy data item X*. When Runnable R1 is running in higher priority task context than R2, and R1 is the only one writing X and R2 is reading X it is possible to guarantee data consistency by making a copy of data item X to variable X* before activation of R2 and redirecting write access from X to X* or the read access from X to X* for R2.

- Application for **write conflicts**:

If one or more data item receiver with a higher priority than the sender exist, a *write copy* for the sender is provided.

Example:

There exist Runnable R1, Runnable R2, data item X and copy data item X*. When Runnable R1 (running in lower priority task context than R2) is writing X and R2 is reading X, it is possible to guarantee data consistency by making a copy of data item X to data item X* before activation of R1 together with redirecting the write access from X to X* for R1 or the read access from X to X* for R2.

Usage of this copy mechanism may make sense if one or more of the following conditions hold:

- This copy mechanism can handle those cases when only one instance does the data write access.
- R2 is accessing X several times.
- More than one Runnable R2 has read (resp. write) access to X.
- To save runtime is more important than to save code and RAM.
- Additional RAM requirements to hold the copies is acceptable.

Further issues to be taken into account:

- AUTOSAR SW-Cs provided as source code and AUTOSAR SW-Cs provided as object code may or have to be handled in different ways. The redirecting mechanism for source code could use macros for C and C++ very efficiently whereas object-code AUTOSAR SW-Cs most likely are

forced to use references.

Note that the copy strategy is used to guarantee data consistency for implicit sender-receiver communication (realizing `DataReadAccess` and `DataWriteAccess`) and for AUTOSAR SW-C internal communication using `InterRunnableVariables` with implicit behavior.

4.2.4.5 Exclusive Areas

The concept of `ExclusiveArea` is more a working model. It's not a concrete implementation approach, although concrete possible mechanisms are listed in AUTOSAR SW-C template specification [17].

Focus of the `ExclusiveArea` concept is to block potential concurrent accesses to get data consistency.

`ExclusiveAreas` are associated with `Runnable`s. The RTE is forced to guarantee data consistency when the `Runnable` runs in an `ExclusiveArea`. A `Runnable` can run inside one or several `ExclusiveAreas` completely or can enter one or several `ExclusiveAreas` during their execution for one or several times .

- If an AUTOSAR SW-C requests the RTE to look for data consistency for it's internally used data (for a part of it or the complete one) using the `ExclusiveArea` concept, the SW designer can use the API calls "`Rte_Enter()`" in 5.6.24 and "`Rte_Exit()`" in 5.6.25 to specify where he wants to have the protection by RTE applied.
"`Rte_Enter()`" defines the begin and "`Rte_Exit()`" defines the end of the code sequence containing data accesses the RTE has to guarantee data consistency for.
- If the SW designer wants to have the mutual exclusion for complete `Runnable`s he can specify this by setting the attribute "`RunnableEntityRunsInExclusiveArea`" in the AUTOSAR SW-C description.

In principle the `ExclusiveArea` concept can handle the access to single data items as well as the access to several data items realized by a group of instructions. It also doesn't matter if one `Runnable` is completely running in an `ExclusiveArea` and another `Runnable` only temporarily enters the same `ExclusiveArea`. The RTE has to guarantee data consistency.

[rte_sws_3500] The RTE has to guarantee data consistency for arbitrary accesses to data items accessed by `Runnable`s marked with the same `ExclusiveArea`.

[rte_sws_3515] RTE has to provide an API enabling the SW-Cs to access and leave `ExclusiveAreas`.

If `Runnable`s accessing same `ExclusiveArea` are assigned to be executing in different task contexts, the RTE can apply suitable mechanisms, e.g. task blocking, to guarantee

data consistency for data accesses in the common ExclusiveArea. However, special attributes can be set that require certain data consistency mechanisms in which case the RTE generator is forced to apply the selected mechanism.

4.2.4.5.1 Assignment of data consistency mechanisms

There might be domain, ECU or even project specific needs which data consistency mechanism makes sense most to be applied to an ExclusiveArea. The decision which mechanism has to be applied by RTE is taken during ECU integration by setting the ExclusiveArea configuration parameter *ExclusiveAreaImplMechanism*. This parameter is an input for RTE generator.

As stated in section 4.2.4.4 there might be more mechanisms to realize ExclusiveAreas as mentioned in this specification. So RTE implementations might provide other mechanisms in plus by a vendor specific solutions. This allows further optimizations.

Actually following values for configuration parameter *ExclusiveAreaImplMechanism* must be supported:

- **InterruptBlocking**
This value requests enabling and disabling of Interrupts and is based on the *Interrupt blocking strategy*.
- **OSResources**
This value requests to apply the *Usage of OS resources* mechanism.
- **NonPreemptiveTasks**
This value requests to apply the *Task blocking strategy*.
- **CooperativeRunnablePlacement**
This value requires to apply the *Cooperative Runnable Placement Strategy*.

The strategies / mechanisms are described in general in section 4.2.4.4.

[rte_sws_3504] If the configuration parameter *ExclusiveAreaImplMechanism* of an ExclusiveArea is set to value *"InterruptBlocking"* the RTE generator shall use the mechanism of *Interrupt blocking* to guarantee data consistency if data inconsistency could occur.

[rte_sws_3595] If the configuration parameter *ExclusiveAreaImplMechanism* of an ExclusiveArea is set to value *"OSResources"* the RTE generator shall use OS resources to guarantee data consistency if data inconsistency could occur.

The requirements above have the limitation "if data inconsistency could occur" because it makes no sense to apply a data consistency mechanism if no potential data inconsistency can occur. This can be relevant if e.g. the "Sequential scheduling strategy" (described in section 4.2.4.4) still has solved the item by the ECU integrator

defining an appropriate runnable-to-task mapping.

[rte_sws_3596] If the configuration parameter `ExclusiveAreaImplMechanism` of an `ExclusiveArea` is set to value *"NonPreemptiveTasks"* the RTE generator shall generate code relying on the fact that the task containing `Runnable`s accessing the `ExclusiveArea` cannot be preempted by other tasks containing `Runnable`s accessing the same `ExclusiveArea` too.

For usage of value *"NonPreemptiveTasks"* it is assumed that the corresponding runnable to task mapping and the task configuration has been done appropriate.

[rte_sws_3503] If the configuration parameter `ExclusiveAreaImplMechanism` of an `ExclusiveArea` is set to value *"CooperativeRunnablePlacement"* the RTE generator shall generate code according the *Cooperative Runnable Placement Strategy* to guarantee data consistency.

Since the decision to select the `Cooperative` runnable placement strategy to prohibit data access conflicts affects the behavior of several tasks and potentially many `ExclusiveAreas` the RTE generator is not allowed to override the decision.

In a SWC code, it is not allowed to use waitpoints inside an `ExclusiveArea`: The RTE generator might use OSEK services to implement `ExclusiveAreas` and waiting for an OS event is not allowed when an OSEK resource has been taken for example. For `RunnableEntityEntersExclusiveArea`, the RTE generator cannot check if waitpoints are inside an `ExclusiveArea`. Therefore, this is the responsibility of the SWC Code writer to ensure that no wait points are used inside exclusive area. But for runnable *"RunnableEntityRunsInExclusiveArea"*, the RTE generator is able to do the following check.

[rte_sws_7005] The RTE generator shall reject a configuration with a waitpoint applied to a runnable for which the attribute *"RunnableEntityRunsInExclusiveArea"* is set.

4.2.4.6 InterRunnableVariables

A non-composite AUTOSAR SW-C can reserve `InterRunnableVariables` which can be accessed by the `Runnable`s of this one AUTOSAR SW-C (also see section 4.3.3.1). Read and write accesses are possible. There is a separate set of those variables per AUTOSAR SW-C instance.

Again the RTE has to guarantee data consistency. Appropriate means will depend on `Runnable` placement decisions which are taken during ECU configuration.

[rte_sws_3516] The RTE has to guarantee data consistency for communication between `Runnable`s of one AUTOSAR SW-Component instance using the same `InterRunnableVariable`.

Next the two kinds of `InterRunnableVariables` are treated:

1. InterRunnableVariables with **implicit** behavior
2. InterRunnableVariables with **explicit** behavior

4.2.4.6.1 InterRunnableVariables with implicit behavior

In applications with very high SW-C communication needs and much real time constraints (like in powertrain domain) the usage of a copy mechanism to get data consistency might be a good choice because during Runnable execution no data consistency overhead in form of concurrent access blocking code and runtime during its execution exists - independent of the number of data item accesses.

Costs are code overhead in the Runnable prolog and epilog which is often be minimal compared to other solutions. Additional RAM need for the copies comes in plus.

When *InterRunnableVariables with implicit behavior* are used the RTE is required to make the data available to the Runnable using the semantic of a copy operation but is not necessarily required to use a unique copy for each Runnable.

Focus of *InterRunnableVariable with implicit behavior* is to avoid concurrent accesses by redirecting second, third, .. accesses to data item copies.

[rte_sws_3517] The RTE shall guarantee data consistency for *InterRunnableVariables with implicit behavior* by avoiding concurrent accesses to data items specified by InterRunnableVariables using one or more copies and redirecting accesses to the copies.

Compared with Sender/Receiver communication

- Like with DataReadAccess/DataWriteAccess the Runnable IN data is stable during Runnable execution, which means that during an Runnable execution several read accesses to an InterRunnableVariable always deliver the same data item value.
- Like with DataWriteAccess/DataWriteAccess the Runnable OUT data is forwarded to other Runnables not before Runnable execution has terminated, which means that during an Runnable execution write accesses to InterRunnableVariable are not visible to other Runnables.

This behavior requires that Runnable execution terminates.

[rte_sws_3582] Several read accesses to *InterRunnableVariables with implicit behavior* during a Runnable execution shall always deliver the same data item value.

[rte_sws_3583] Several write accesses to *InterRunnableVariables with implicit behavior* during a Runnable execution shall result in only one update of the InterRunnableVariable content visible to other Runnables with the last written value.

[rte_sws_3584] The update of *InterRunnableVariables with implicit behavior* done during a Runnable execution shall be made available to other Runnables after the Runnable execution has terminated.

The usage of *InterRunnableVariables with implicit behavior* shall be valid for category 1a and 1b Runnable entities. Usage in category 2 (and 3) Runnables is not allowed because there Runnable termination is not guaranteed and so it's not guaranteed that other Runnables will ever get the updated data. See also requirement rte_sws_3518.

For API of *InterRunnableVariables with implicit behavior* see sections 5.6.20 and 5.6.21.

For more details how this mechanism could work see "Copy strategy" in section 4.2.4.4.

4.2.4.6.2 InterRunnableVariables with explicit behavior

In many applications saving RAM is more important than saving runtime. Also some application require to have access to the newest data item value without any delay, even several times during execution of a Runnable.

Both requirements can be fulfilled when RTE supports data consistency by blocking second/third/.. concurrent accesses to a signal buffer if data consistency is jeopardized. (Most likely RTE has nothing to do if SW is running on a 16bit machine and making an access to an 16bit value when a 16bit data bus is present.)

Focus of *InterRunnableVariables with explicit behavior* is to block potential concurrent accesses to get data consistency.

The mechanism behind is the same as in the ExclusiveArea concept (see section 4.2.4.5). But although ExclusiveAreas can handle single data item accesses too, their API is made to make the RTE to apply data consistency means for a group of instructions accessing several data items as well. So when using an ExclusiveArea to protect accesses to one single common used data item each time two RTE API calls grouped around are needed. This is very inconvenient and might lead to faults if the calls grouped around might be forgotten.

The solution is to support *InterRunnableVariables with explicit behavior*.

[rte_sws_3519] The RTE shall guarantee data consistency for *InterRunnableVariables with explicit behavior* by blocking concurrent accesses to data items specified by InterRunnableVariables.

The RTE generator is not free to select on it's own if implicit or explicit behavior shall be applied. Behavior must be known at AUTOSAR SW-C design time because in case of *InterRunnableVariables with implicit behavior* the AUTOSAR SW-C designer might rely on the fact that several read accesses always deliver same data item value.

[rte_sws_3580] The RTE shall supply different APIs for *InterRunnableVariables with implicit* behavior and *InterRunnableVariables with explicit* behavior.

For API of *InterRunnableVariables with explicit behavior* see sections 5.6.22 and 5.6.23.

4.2.5 Multiple trigger of Runnables

Concurrent activation

The AUTOSAR SW-C template specification [17] states that runnable entities (further called "Runnables") might be invoked concurrently several times if the Runnables attribute "canBeInvokedConcurrently" is set. It's then in the responsibility of the AUTOSAR SW-C designer that no data might be corrupted when the Runnable is activated several times in parallel.

[rte_sws_3523] The RTE has to support concurrent activation of the same instance of a runnable entity if the associative attribute "canBeInvokedConcurrently" is set to TRUE. This includes concurrent activation in several tasks. If the attribute is not set resp. set to FALSE, concurrent activation of the runnable entity is forbidden.

Activation by several RTEEvents

Nevertheless a Runnable whose attribute "canBeInvokedConcurrently" is NOT set might be still activated by several RTEEvents if activation configuration guarantees that concurrent activation can never occur. This includes activation in different tasks. A standard use case is the activation of same instance of a runnable in different modes.

[rte_sws_3520] The RTE supports activation of same instance of a runnable entity by multiple RTEEvents.

RTEEvents are triggering Runnable activation and may supply 0..several role parameters, see *section 5.7.3*. Role parameters are not visible in the Runnables signature - except in those triggered by an *OperationInvokedEvent*. With the exception of the RTEEvent *OperationInvokedEvent* all role parameters can be accessed by user with implicit or explicit Receiver API.

[rte_sws_3524] The RTE supports activation of same instance of a runnable entity by RTEEvents of different kinds.

The RTE shall NOT support a runnable entity triggered by an RTEEvent *OperationInvokedEvent* to be triggered by any other RTEEvent except for other *OperationInvokedEvents* of compatible operations. This limitation is stated in appendix in *section A.2*.

4.2.6 Measurement and Calibration

4.2.6.1 General

Calibration is the process of adjusting an ECU SW to fulfill its tasks to control physical processes respectively to fit it to special project needs or environments. To do this two different mechanisms are required and have to be distinguished:

1. Measurement

Measure what's going on in the ECU e.g. by monitoring communication data (Inter-ECU, Intra-ECU, Intra-SWC). There are several ways to get the monitor data out of the ECU onto external visualization and interpretation tools.

2. Calibration

Based on the measurement data the ECU behavior is modified by changing parameters like runtime SW switches, process controlling data of primitive or complex type, interpolation curves or interpolation fields. In the following for such parameters the term calibration parameter is used.

Online and offline calibration

The way how measurement and calibration is performed is company, domain and project specific. Nevertheless two different basic situations can be distinguished and are important for understanding:

1. Offline calibration

Measure when ECU is running, change calibration data when ECU is off.
Process might look like this:

- (a) Flash the ECU with current program file
- (b) PowerUp ECU in target (actual or emulated) environment
- (c) Measure running ECU behavior - log or monitor via external tooling
- (d) Switch off ECU
- (e) Change calibration parameters and create a new flashable program file (hex-file) e.g. by performing a new SW make run
- (f) Back to (a).

Do loop as long as a need for calibration parameter change exists or the Flash survives.

2. Online calibration

Do measurement and calibration in parallel.

In this case in principle all steps mentioned in "Offline calibration" above have to be performed in parallel. So other mechanisms are introduced avoiding ECU flashing when modifying ECU parameters. ECU works temporarily with changed

data and when the calibration process is over the result is an updated set of calibration data. In next step this new data set might be merged into the existing program file or the new data set might be an input for a new SW make run. In both cases the output is a new program file to flash into the ECU.

Process might look like this:

- (a) Flash the ECU with current program file
- (b) PowerUp ECU in target environment
- (c) Measure running ECU behavior and temporarily modify calibration parameters. Store set of updated calibration parameters (not on the ECU but on the calibration tool computer). Actions in step c) may be done iteratively.
- (d) Switch off ECU
- (e) Create a new flashable program file (hex-file) containing the new calibration parameters

Procedure over

4.2.6.2 Measurement

4.2.6.2.1 What can be measured

The AUTOSAR SW-C template specification [17] explains to which AUTOSAR prototypes a measurement pattern can be applied.

RTE provides measurement support for

1. communication between Ports

Measurable are

- DataElementPrototypes of a SenderReceiverInterface used in a PortPrototype (of a ComponentPrototype) to capture sender-receiver communication between ComponentPrototypes
- ArgumentPrototypes of an OperationPrototype in a ClientServerInterface to capture client-server communication between ComponentPrototypes

2. communication inside of AUTOSAR SW-Cs

Measurable are InterrunnableVariables

4.2.6.2.2 RTE support for Measurement

The way how measurement data is read out of the ECU is not focus of the RTE specification. But the RTE structure and behavior must be specified in that way that measurement values can be provided by RTE during ECU program execution.

To avoid synchronization effort it shall be possible to read out measurement data asynchronously to RTE code execution. For this the measurement data must be stable. As a consequence this might forbid direct reuse of RAM locations for implementation of several AUTOSAR communications which are independent of each other but occurring sequentially in time (e.g. usage of same RAM cell to store UInt8 data sender receiver communication data between Runnables at positions 3 and 7 and later the same RAM cell for the communication between Runnables at positions 9 and 14 of same periodically triggered task). So applying measurable elements might lead to less optimizations in the generated RTE's code and to increased RAM need.

There are circumstances when RTE will store same communication data in different RAM locations, e.g. when realizing implicit sender receiver communication or InterRunnableVariables with implicit behavior. In these cases there is only the need to have the content of one of these stores made accessible from outside.

The information that measurement shall be supported by RTE is defined in applied SwDataDefProps:

The value *READ-ONLY* of the property *swCalibrationAccess* defines that measurement shall be supported, any other value of the property *swCalibrationAccess* is to be ignored for measurement.

Following requirements *rte_sws_3900*, *rte_sws_3901* and *rte_sws_3902* cover 2 cases each where SwDataDefProps can be applied:

- On level of type (only supported for primitive types):
Every instance of the type is to be measured
- On level of *DataPrototype* and its specializations (e.g. *ArgumentPrototype* or *InterrunnableVariable*) (supported for all types):
The settings override the settings given by the type

SwDataDefProps contain more informations how measurement values are to be interpreted and presented by external calibration tools. These informations are needed for the ASAM2 respectively A2L file generation. Afterwards the A2L file is used by ECU-external measurement and calibration tools so that these tools know e.g. how to interpret raw data received from ECU and how to get them.

SwDataDefProps also contain information about which section shall be used to allocate memory for measurement data. For that see section 4.2.6.3.5.7.

For sender-receiver resp. client-server communication same or compatible interfaces are used to specified connected ports. So very often measurement will be demanded two times for same or compatible DataElementPrototype on provide and require side of a 1:1 communication resp. multiple times in case of 1:N or M:1 communication. In that case providing more than one measurement value for a DataElementPrototype doesn't make sense and would increase ECU resources need excessively. Instead only one measurement value shall be provided.

Sender-receiver communication

[rte_sws_3900] If the property *swCalibrationAccess* enclosed in the *SwDataDefProps* of a *DataPrototype* or its associated type used in an interface of a sender-receiver port of a *ComponentPrototype* is set to *READ-ONLY* the RTE generator has to provide one reference to a location in memory where the actual content of the instance specific data of the corresponding *DataElementPrototype* of the communication can be read.

To prohibit multiple measurement values for same communication:

(Note that affected *DataElementPrototypes* might be specified in same or compatible port interfaces.)

[rte_sws_3972] For 1:1 and 1:N sender-receiver communication the RTE shall provide measurement values taken from sender side if measurement is demanded in provide and require port.

[rte_sws_3973] For N:1 intra-ECU sender-receiver communication the RTE shall provide measurement values taken from receiver side if measurement is demanded in provide and require ports.

Note:

See further below for support of queued communication.

[rte_sws_3974] For a *DataElementPrototype* with measurement demand associated with received data of inter-ECU sender-receiver communication the RTE shall provide only one measurement store reference containing the actual received data even if several receiver ports demand measurement.

Client-Server communication

[rte_sws_3901] If the property *swCalibrationAccess* enclosed in the *SwDataDefProps* of an *ArgumentPrototype* or its associated type used in an interface of a client-server port of a *ComponentPrototype* is set to *READ-ONLY* the RTE generator has to provide one reference to a location in memory where the actual content of the instance specific argument data of the communication can be read.

To prohibit multiple measurement values for same communication:

(Note that affected *DataElementPrototypes* might be specified in same or compatible port interfaces.)

[rte_sws_3975] For intra-ECU client-server communication the RTE shall provide measurement values taken from client side if measurement of an *ArgumentPrototype* is demanded by provide and require ports.

[rte_sws_3976] For inter-ECU client-server communication with the client being present on same ECU as the RTE, the RTE shall provide measurement values taken from client side.

[rte_sws_3977] For inter-ECU client-server communication with the server being present on same ECU as the RTE, the RTE shall provide measurement values taken

from server if no client present on same ECU as the server is connected with that server too.

Note:

When a measurement is applied to a client-server call additional copy code might be produced so that a zero overhead direct server invocation is no longer possible for this call.

InterrunnableVariables

[rte_sws_3902] If the property *swCalibrationAccess* enclosed in the *SwDataDefProps* of an *InterRunnableVariable* or its associated type is set to *READ-ONLY* the RTE generator has to provide one reference to a location in memory where the actual content of the *InterRunnableVariable* can be read for a specific instantiation of the AUTOSAR SWC.

Unconnected ports or compatible interfaces

As stated in section 5.2.7 RTE supports handling of unconnected ports.

Measurement support for unconnected sender-receiver provide ports makes sense since a port might be intentionally added for monitoring purposes only.

Measurement support for unconnected sender-receiver require ports makes sense since the measurement is specified on the type level of the Software Component and therefore independent of the individual usage of the Software Component. In case of unconnected sender-receiver require ports the measurement shall return the initial value.

Support for unconnected client-server provide port does not make sense since the server cannot be called and with this no data can be passed there.

Support for unconnected client-server require port makes sense since the measurement is specified on the type level of the Software Component and therefore independent of the individual usage of the Software Component. In case of unconnected client-server require ports the measurement shall return the actually provided and returned values.

[rte_sws_3978] For sender-receiver communication the RTE generator shall respect measurement demands enclosed in unconnected provide ports.

[rte_sws_5101] For sender-receiver communication the RTE generator shall respect measurement demands enclosed in unconnected require ports and deliver the initial value.

[rte_sws_3980] For client-server communication the RTE generator shall ignore measurement demands enclosed in unconnected provide ports.

[rte_sws_5102] For client-server communication the RTE generator shall respect measurement demands enclosed in unconnected require ports.

Principally same thoughts as above are applied to unused DataElementPrototypes for sender-receiver communication where ports with compatible but not same interfaces are connected. It's no issue for client-server due to compatibility rules for client-server interfaces since in compatible client-server interfaces all OperationPrototypes have to be present in provide and require port (see AUTOSAR SW-C Template [17]).

[rte_sws_3979] For sender-receiver communication the RTE generator shall respect measurement demands of those DataElementPrototypes in connected ports when provide and require port interfaces are not the same (but only compatible) even when a DataElementPrototype in the provide port has no assigned DataElementPrototype in the require port.

General measurement disabling switch

To support saving of ECU resources for projects where measurement isn't required at all whereas enclosed AUTOSAR SW-Cs contain SwDataDefProps requiring it, it shall be possible to switch off support for measurement. This shall not influence support for calibration (see 4.2.6.3).

[rte_sws_3903] The RTE generator shall have the option to switch off support for measurement for generated RTE code. This option shall influence complete RTE code at once.

There also might be projects in which monitoring of ECU internal behavior is required but calibration is not.

[rte_sws_3904] The enabling of RTE support for measurement shall be independent of the enabling of the RTE support for calibration.

Queued communication

Measurement of queued communication is not supported yet. Reasons are:

- A queue can be empty. What's to measure then?
- Which of the queue entries is the one to take the data from might differ out of user view?
- Only quite inefficient solutions possible because implementation of queues entails storage of information dynamically at different memory locations. So always additional copies are required.

[rte_sws_3950] RTE generator shall reject configurations where measurement for queued communication is configured.

4.2.6.3 Calibration

The RTE has to support the allocation of calibration parameters and the access to them for SW using them. As seen later on for some calibration methods the RTE must contain support SW too (see 4.2.6.3.5).

But in general the RTE is not responsible for the exchange of the calibration data values or the transportation of them between the ECU and external calibration tools.

4.2.6.3.1 Calibration parameters

Calibration parameters (which the AUTOSAR SW-C template specification [17] calls `CalprmElements`) can be defined in `CalprmComponentTypes` and in AUTOSAR SW-Cs.

1. `CalprmComponentTypes` don't have an internal behavior but contain `CalprmElementPrototypes` and serve to provide calibration parameters used commonly by several AUTOSAR SW-Cs. The use case that one or several of the user SW-Cs are instantiated on different ECUs is supported by instantiation of the `CalprmComponentType` on the affected ECUs too.

Of course several AUTOSAR SW-Cs allocated on one ECU can commonly access the calibration parameters of `CalprmComponentTypes` too. Also several instances of an AUTOSAR SW-Cs can share the same calibration parameters of a `CalprmComponentType`.

2. Calibration parameters defined in AUTOSAR SW-Cs can only be used inside the SW-C and are not visible to other SW-Cs. Instance individual and common calibration parameters accessible by all instances of a AUTOSAR SW-C are possible.

[rte_sws_3958] Several AUTOSAR SW-Cs (and also several instances of AUTOSAR SW-Cs) shall be able to share same calibration parameters defined in `CalprmComponentTypes`.

[rte_sws_7186] The generated RTE shall initialize the memory objects implementing `CalprmElementPrototypes` in *p-ports* of `CalprmComponentTypes` according to the *ValueSpecification* of the `ParameterProvideComSpec` referring the `CalprmElementPrototype` in the *p-port*.

[rte_sws_3959] If the attribute "perInstanceCalprm" of a `CalprmElementPrototype` of an ATOMIC SW-C is set to TRUE the RTE shall support the access to instance specific calibration parameters of the AUTOSAR SW-C. If the attribute of a `CalprmElementPrototype` of an ATOMIC SW-C is set to FALSE access to a the RTE shall create a common access to the shared calibration parameter.

[rte_sws_7185] The generated RTE shall initialize the memory objects implementing `CalprmElementPrototypes` in the role `perInstanceCalprm` or `sharedCalprm` if a `LocalParameterInitValueAssignment` (part of

InternalBehavior) exists referring to the CalprmElementPrototype according to this ValueSpecification.

It might be project specific or even project phase specific which calibration parameters have to be calibrated and which are assumed to be stable. So it shall be selectable on CalprmComponentTypes and AUTOSAR SW-C granularity level for which calibration parameters RTE shall support calibration.

[rte_sws_3905] It shall be configurable for each CalprmComponentType if RTE calibration support for the enclosed CalprmElementPrototypes is enabled or not.

[rte_sws_3906] It shall be configurable for each AUTOSAR SW-C if RTE calibration support for the enclosed CalprmElementPrototypes is enabled or not.

RTE calibration support means the creation of SW as specified in section 4.2.6.3.5 "Data emulation with SW support".

Require ports on CalprmComponentTypes don't make sense. CalprmComponentTypes only have to provide calibration parameters to other Component types. So the RTE generator shall reject configurations containing require ports attached to CalprmComponentTypes. (see section A.14)

4.2.6.3.1.1 Separation of calibration parameters

Sometimes it is required that one or more calibration parameters out of the mass of calibration parameters of an CalprmComponentType respectively an AUTOSAR SW-C shall be placed in another memory location than the other parameters of the CalprmComponentType respectively the AUTOSAR SW-C. This might be due to security reasons (separate normal operation from monitoring calibration data in memory) or the possibility to change calibration data during a diagnosis session (which the calibration parameter located in NVRAM).

[rte_sws_3907] The RTE generator shall support separation of calibration parameters from CalprmComponentTypes respectively AUTOSAR SW-Cs depending on the CalprmElementPrototype property "swAddrMethod".

4.2.6.3.2 Support for offline calibration

As described in section 4.2.6.1 when using an offline calibration process measurement is decoupled from providing new calibration parameters to the ECUs SW. During measurement phase information is collected needed to define to which values the calibration parameters are to be set best. Afterwards the new calibration parameter set is brought into the ECU e.g. by using a bootloader.

[rte_sws_3971] The RTE generator shall have the option to switch off all *data emulation* support for generated RTE code. This option shall influence complete RTE code at once.

The term *data emulation* is related to mechanisms described in section 4.2.6.3.3.

Out of view of RTE the situation is same as when *data emulation without SW support* (described in section 4.2.6.3.4) is used:

The RTE is only responsible to provide access to the calibration parameters via the RTE API as specified in section 5.6. Exchange of CalprmElement content is done invisibly for ECU program flow and with this for RTE too.

When no *data emulation support* is required calibration parameter accesses to parameters stored in FLASH could be performed by direct memory read accesses without any indirection for those cases when accesses are coming out of single instantiated AUTOSAR SW-Cs. Nevertheless it's not goal of this specification to require direct accesses since this touches implementation. It might be ECU HW dependent or even be project dependent if other accesses are more efficient or provide other significant advantages or not.

4.2.6.3.3 Support for online calibration: Data emulation

To allow **online calibration** it must be possible to provide alternative calibration parameters invisible for application. The mechanisms behind are described here. We talk of *data emulation*.

In the following several calibration methods are described:

1. Data emulation without SW support and
2. several methods of data emulation with SW-support.

The term **data emulation** is used because the change of calibration parameters is emulated for the ECU SW which uses the calibration data. This change is invisible for the user-SW in the ECU.

RTE is significantly involved when SW support is required and has to create calibration method specific SW. Different calibration methods means different support in Basic SW which typically is ECU integrator specific. So it does not make sense to support DIFFERENT data emulation with SW support methods in ANY one RTE build. But it makes sense that the RTE supports direct access (see section 4.2.6.3.4) for some AUTOSAR SW-Cs resp. CalprmComponentTypes and one of the data emulation with SW support methods (see section 4.2.6.3.5) for all the other AUTOSAR SW-Cs resp. CalprmComponentTypes at the same time.

[rte_sws_3909] The RTE shall support only one of the data emulation with SW support methods at once.

4.2.6.3.4 Data emulation without SW support (direct access)

For "online calibration" (see section 4.2.6.1) the ECU is provided with additional hardware which consists of control logic and memory to store modified calibration parameters in. During ECU execution the brought in control logic redirects memory accesses to new brought in memory whose content is modified by external tooling without disturbing normal ECU program flow. Some microcontrollers contain features supporting this. A lot of smaller microcontrollers don't. So this method is highly HW dependent.

To support these cases the RTE doesn't have to provide e.g. a reference table like described in section 4.2.6.3.5. Exchange of `CalprmElement` content is done invisibly for program flow and for RTE too.

[rte_sws_3942] The RTE generator shall have the option to switch off *data emulation with SW support* for generated RTE code. This option shall influence complete RTE code at once.

4.2.6.3.5 Data emulation with SW support

In case "online calibration" (see section 4.2.6.1) is required, quite often data emulation without support by special SW constructs isn't possible. Several methods exist, all have the consequence that additional need of ECU resources like RAM, ROM/FLASH and runtime is required.

Data emulation with SW support is possible in different manners. During calibration process in each of these methods modified calibration data values are kept typically in RAM. Modification is controlled by ECU external tooling and supported by ECU internal SW located in AUTOSAR basic SW or in complex driver.

If calibration process isn't active the accessed calibration data is originated in ROM/FLASH respectively in NVRAM in special circumstances (as seen later on).

Since multiple instantiation is to be supported several instances of the same `CalprmElementPrototypes` have to be allocated. Because the RTE is the only one SW in an AUTOSAR ECU able to handle the different instances the access to these calibration parameters can only be handled by the RTE. So the RTE has to provide additional SW constructs required for data emulation with SW support for calibration.

However the RTE doesn't know which of the ECU functionality shall be calibrated during a calibration session. To allow expensive RAM to be reused to calibrate different ECU functionalities in one or several online calibration sessions (see 4.2.6.1) in case of the single and double pointer methods for data emulation with SW support described below the RTE has only to provide the access to `CalprmElements` during runtime but allowing other SW (a BSW module or a complex driver) to redirect the access to alternative calibration parameter values (e.g. located in RAM) invisibly for application.

The RTE is neither the instance to supply the alternative values for CalprmElements nor in case of the pointered methods for data emulation with SW support to do the redirection to the alternative values.

[rte_sws_3910] The RTE shall support *data emulation with SW support* for calibration.

[rte_sws_3943] The RTE shall support these data emulation methods with SW support:

- Single pointered calibration parameter access further called "single pointered method"
- Double pointered calibration parameter access further called "double pointered method"
- Initialized RAM parameters further called "initRAM parameter method"

To save RAM/ROM/FLASH resources in single pointered method and double pointered method CalprmElement allocation is done in groups. One entry of the calibration reference table references the begin of a group of CalprmElements. For better understanding successional this group is called **CalprmElementGroup** (which is no term out of the AUTOSAR SW-C template specification [17]). One CalprmElementGroup can contain one or several CalprmElements.

[rte_sws_3911] If data emulation with SW support is enabled, the RTE generator shall allocate all CalprmElements marked with same property "swAddrMethod" of one instance of a CalprmComponentType consecutively. Together they build a separate CalprmElementGroup.

[rte_sws_3912] If data emulation with SW support is enabled, the RTE shall guarantee that all CalprmElements marked with same property "swAddrMethod" of an AUTOSAR SWC instance are allocated consecutively. Together they build a separate CalprmElementGroup.

It is not possible to access same calibration parameter inside of a CalprmComponentType via several ports. This is a consequence of the need to support the use case that a CalprmComponentType shall be able to contain several calibration parameters derived from one CalprmElementPrototype which is contained in one interface applied to several ports of the CalprmComponentType. Using only the CalprmElementPrototype names for the names of the elements of a CalprmElementGroup would lead to a name clash since then several elements with same name would have to be created. So port prototype and CalprmElementPrototype name are concatenated to specify the CalprmElementGroup member names.

This use case cannot be applied to AUTOSAR SW-C internal calibration parameters since they cannot be accessed via AUTOSAR ports.

[rte_sws_3968] The names of the elements of a CalprmElementGroup derived from a CalprmComponentType shall be <port>_<element> where <port> is the short-

name of the provided AUTOSAR port prototype and <element> the short-name of the CalprmElementPrototype within the CalPrmInterface categorizing the PPort.

4.2.6.3.5.1 Single pointered method

There is one calibration reference table in RAM with references to one or several CalprmElementGroups. Accesses to calibration parameters are indirectly performed via this reference table.

Action during calibration procedure e.g. calibration parameter value exchange is not focus of this specification. Nevertheless an example is given for better understanding.

Example how the exchange of calibration parameters could be done for single pointered method:

1. Fill a RAM buffer with the modified calibration parameter values for complete CalprmElementGroup
2. Modify the corresponding entry in the calibration reference table so that a redirection to new CalprmElementGroup is setup

Now calibration parameter accesses deliver the modified values.

Figure figure 4.21 illustrates the method.

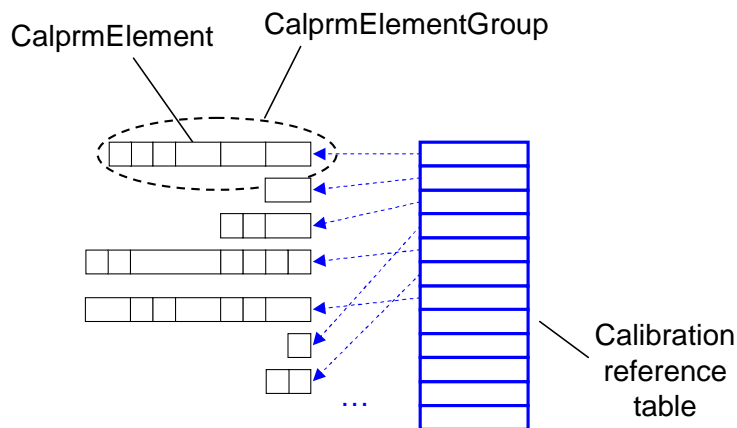


Figure 4.21: CalprmElementGroup in single pointered method context

[rte_sws_3913] If data emulation with SW support with single pointered method is enabled, the RTE generator shall create a table located in RAM with references to CalprmElementGroups. The type of the table is of ArrayType.

One reason why in this approach the calibration reference table is realized as an array is to make ECU internal reference allocation traceable for external tooling. Another is to allow a Basic-SW respectively a complex driver to emulate other calibration parameters which requires the standardization of the calibration reference table too.

[rte_sws_3947] If data emulation with SW support with single method is enabled the name (the label) of the calibration reference table shall be `<RteCalprmRefTab>`.

Calibration parameters located in NVRAM are handled same way (also see section 4.2.6.3.6).

[rte_sws_3936] If data emulation with SW support with single or double pointered method is enabled and calibration parameter respectively a `CalprmElementGroups` is located in NVRAM the corresponding calibration reference table entry shall reference the `PerInstanceMemory` working as the NVRAM RAM buffer.

4.2.6.3.5.2 Double pointered method

There is one calibration reference table in ROM respectively Flash with references to one or several `CalprmElementGroups`. Accesses to calibration parameters are performed through a double indirection access. During system startup the base reference is initially filled with a reference to the calibration reference table.

Action during calibration procedure e.g. calibration parameter value exchange is not focus of this specification. Nevertheless an example is given for better understanding.

Example how the exchange of calibration parameters could be done for double pointered method:

1. Copy the calibration reference table into RAM
2. Fill a RAM buffer with modified calibration parameter values for complete `CalprmElementGroup`
3. Modify the corresponding entry in the RAM copy of the reference table so that a redirection to new `CalprmElementGroup` is setup
4. Change the content of the base reference so that it references the calibration reference table copy in RAM.

Now calibration parameter accesses deliver the modified values.

[rte_sws_3914] If data emulation with SW support with double pointered method is enabled, the RTE generator shall create a table located in ROM respectively FLASH with references to `CalprmElementGroups`. The type of the table is of `ArrayType`.

Figure figure 4.22 illustrates the method.

To allow a Basic-SW respectively a complex driver to emulate other calibration parameters the standardization of the base reference is required.

[rte_sws_3948] If data emulation with SW support with double method is enabled the name (the label) of the calibration base reference shall be `<RteCalprmBase>`. This label and the base reference type shall be exported and made available to other SW

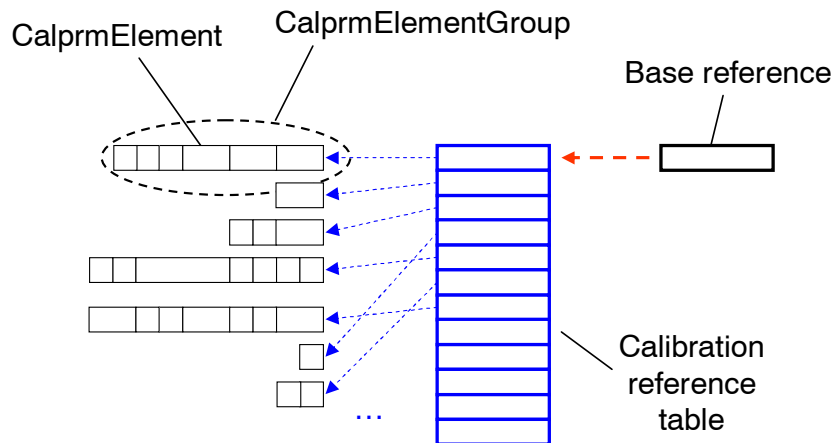


Figure 4.22: CalprmElementGroup in double pointered method context

on same ECU.

Calibration parameters located in NVRAM are handled same way (also see section 4.2.6.3.6).

For handling of calibration parameters located in NVRAM with single or double pointered method see `rte_sws_3936` in section 4.2.6.3.5.1. General information is found in section 4.2.6.3.6).

4.2.6.3.5.3 InitRam parameter method

For each instance of a `CalprmElementPrototype` the RTE generator creates a calibration parameter in RAM and a corresponding value in ROM/FLASH. During startup of RTE the calibration parameter values of ROM/FLASH are copied into RAM. Accesses to calibration parameters are performed through a direct access to RAM without any indirection.

Action during calibration procedure e.g. calibration parameter value exchange is not focus of this specification. Nevertheless an example is given for better understanding: An implementation simply would have to exchange the content of the RAM cells during runtime.

[rte_sws_3915] If data emulation with SW support with `initRam` parameter method is enabled, the RTE generator shall create code guaranteeing that

1. calibration parameters are allocated in ROM/Flash and
2. a copy of them is allocated in RAM made available latest during RTE startup

for those `CalprmElementPrototypes` for which calibration support is enabled.




Figure 4.23: initRam Parameter method setup

Figure figure 4.23 illustrates the method.

A special case is the access of CalprmElementPrototypes instantiated in NVRAM (also see section 4.2.6.3.6). In this no extra RAM copy is required because a RAM location containing the calibration parameter value still exists.

[rte_sws_3935] If data emulation with SW support with initRam parameter method is enabled, the RTE generator shall create direct accesses to the PerInstanceMemory working as RAM buffer for the calibration parameters defined to be in NVRAM.

4.2.6.3.5.4 Arrangement of a CalprmElementGroup for pointered methods

For data emulation with SW support with single or double pointered methods the RTE has to guarantee access to each single member of a CalprmElementGroup for source code and object code delivery independent if the member is a primitive or a complex data type. For this the creation of a record type for a CalprmElementGroup was chosen.

[rte_sws_3916] One CalprmElementGroup shall be realized as one record type.

To support object code delivery of CalprmComponents and AUTOSAR SWCs the CalprmElement sequence order in a CalprmElementGroup and in the reference table have to be specified too.

[rte_sws_3917] In compatibility mode the members of a CalprmElementGroup are ordered consecutively according following sequence:

1. At first primitive types:
 - (a) *Double* types (BSW: base type *float64*)
 - (b) *Double_with_NaN* types (BSW: base type *float64*)
 - (c) *Float* types (BSW: REAL-TYPEs with base type *float32*)

- (d) *Float_with_NaN* types (BSW: REAL-TYPEs with base type *float32*)
 - (e) *UInt32* types (BSW: INTEGER-TYPEs with base type *uint32*)
 - (f) *SInt32* types (BSW: INTEGER-TYPEs with base type *sint32*)
 - (g) OPAQUE-TYPEs with base type *uint32*
 - (h) *UInt16* types (BSW: base type *uint16*)
 - (i) *SInt16* types (BSW: base type *sint16*)
 - (j) *Char16* types (BSW: base type *uint16*)
 - (k) OPAQUE-TYPEs with base type *uint16*
 - (l) *UInt8* types (BSW: base type *uint8*)
 - (m) *SInt8* types (BSW: base type *sint8*)
 - (n) *Char8* types (BSW: base type *uint8*)
 - (o) OPAQUE-TYPEs with base type *uint8*
 - (p) *UInt4* types (BSW: base type *uint4*)
 - (q) *SInt4* types (BSW: base type *sint4*)
 - (r) *Boolean* types (BSW: base type *boolean*)
2. Second the complex data types (the elements of AUTOSAR arrays and records are "ordered", therefore there is no need to specify sequence rules):
- (a) array types
 - (b) record types

For AUTOSAR data types also see 5.3.4.

In vendor mode `CalprmElementGroup` order can be redefined to support target specific needs.

[rte_sws_3918] Sequence order of elements of same type in a `CalprmElementGroup` derived from a AUTOSAR SW-C shall be alphabetically (ASCII / ISO 8859-1 code in ascending order) derived from `CalprmElementPrototype` names.

With respect to requirement `rte_sws_3968` this is different for parameters stored in `CalprmComponentTypes`:

[rte_sws_3969] Sequence order of elements of same type in a `CalprmElementGroup` derived from a `CalprmComponentType` shall be alphabetically (ASCII / ISO 8859-1 code in ascending order) at first derived from AUTOSAR port prototype name and if equal secondly derived from `CalprmElementPrototype` names.

4.2.6.3.5.5 Further definitions for pointered methods

As stated in section 4.2.6.3.1.1, dependent of the value of property "swAddrMethod" calibration parameters shall be separated in different memory locations.

[rte_sws_3908] If data emulation with SW support with single or double pointered method is enabled the RTE shall create a separate instance specific CalprmElementGroup for all those CalprmElementPrototypes with a common value of the appended property "swAddrMethod". Those CalprmElementPrototypes which have no property "swAddrMethod" appended, shall be grouped together too.

To allow traceability for external tooling an order must be specified for entries in calibration reference table.

[rte_sws_3920] The entries of the reference table of data emulation with SW support with single or double pointered method shall be a sorted alphabetically (ASCII / ISO 8859-1 code in ascending order) based on the names of the first CalprmElementPrototype member of the referenced CalprmElementGroups.

[rte_sws_3940] Entries in the calibration reference table for data emulation with SW support with single or double pointered method caused by multiple instantiation of CalprmComponentTypes respectively AUTOSAR SW-Cs shall be sorted based on the CalprmComponentTypes respectively AUTOSAR SW-Cs instance names. Sorting rule is ASCII / ISO 8859-1 code in ascending order.

4.2.6.3.5.6 Calibration parameter access

Calibration parameters are derived from CalprmElementPrototypes.

[rte_sws_3921] The RTE has to provide access to each calibration parameter via a separate API call.

API is specified in 5.6.

[rte_sws_3922] If data emulation with SW support and single or double pointered method is enabled the RTE generator shall export the label of the calibration reference table.

[rte_sws_3960] If data emulation with SW support and double pointered method is enabled the RTE generator shall export the label and the type of the calibration base reference.

[rte_sws_3932] If data emulation with SW support with single pointered method is enabled the RTE generator shall create API calls using single indirect access via the calibration reference table for those CalprmElementsPrototypes which are in a CalprmElementGroup for which calibration is enabled.

[rte_sws_3933] If data emulation with SW support with double pointered method is enabled the RTE generator shall create API calls using double indirection access via

the calibration base reference and the calibration reference table for those CalprmElementPrototypes which are in a CalprmElementGroup for which calibration is enabled.

[rte_sws_3934] If data emulation with SW support with double pointered method is enabled, the calibration base reference shall be located in RAM.

4.2.6.3.5.7 Calibration parameter allocation

Since only the RTE knows which instances of AUTOSAR SW-Cs and CalprmComponentTypes are present on the ECU the RTE has to allocate the calibration parameters and reserve memory for them. This approach is also covering multiple instantiated object code integration needs. So memory for instantiated CalprmElementPrototypes is neither provided by CalprmComponentTypes nor by AUTOSAR SW-C.

[rte_sws_3961] The RTE shall allocate the memory for calibration parameters.

A CalprmElementType can be defined to be instance specific or can be shared over all instances of an AUTOSAR SW-C or a CalprmComponentType. The input for the RTE generator contains the values the RTE shall apply to the calibration parameters.

To support online and offline calibration (see section 4.2.6.1) all parameter values for all instances have to be provided.

Background:

- For online calibration often initially the same default values for calibration parameters can be applied. Variation is then handled later by post link tools. Initial ECU startup is not jeopardized. This allows the usage of a default value e.g. by AUTOSAR SW-C or CalprmComponentType supplier for all instances of a CalprmElementPrototype.
- On the other hand applying separate default values for the different instances of a CalprmElementPrototype will be required often for online calibration too, to make a vehicle run initially. This requires additional configuration work e.g. for integrator.
- Offline calibration based on new SW build including new RTE build and compilation process requires all calibration parameter values for all instances to be available for RTE.

Shared CalprmElementPrototypes

[rte_sws_3962] For accesses to a shared CalprmElementPrototype the RTE API shall deliver the same one value independent of the instance the calibration parameter is assigned to.

[rte_sws_3963] The calibration parameter of a shared CalprmElementPrototype shall be stored in one memory location only.

Requirements `rte_sws_3962` and `rte_sws_3963` are to guarantee that only one physical location in memory has to be modified for a change of a shared `CalprmElementPrototype`. Otherwise this could lead to unforeseeable confusion.

Multiple locations are possible for calibration parameters stored in NVRAM. But there a shared `CalprmElementPrototype` is allowed to have only one logical data too.

Instance specific `CalprmElementPrototypes`

[rte_sws_3964] For accesses to an instance specific `CalprmElementPrototype` the RTE API shall deliver a separate calibration parameter value for each instance of a `CalprmElementPrototype`.

[rte_sws_3965] For an instance specific `CalprmElementPrototype` the calibration parameter value of each instance of the `CalprmElementPrototype` shall be stored in a separate memory location.

Usage of `swAddressMethod`

`SwDataDefProps` contain the optional property `swAddressMethod`. It contains meta information about the memory section in which a measurement data store resp. a calibration parameter shall be allocated in. This abstraction is needed to support the reuse of unmodified AUTOSAR SW-Cs resp. `CalprmComponentTypes` in different projects but allowing allocation of measurement data stores resp. calibration parameters in different sections.

Section usage typically depends on availability of HW resources. In one project the micro controller might have less internal RAM than in another project, requiring that most measurement data have to be placed in external RAM. In another project one addressing method (e.g. indexed addressing) might be more efficient for most of the measurement data - but not for all. Or some calibration parameters are accessed less often than others and could be - depending on project specific FLASH availability - placed in FLASH with slower access speed, others in FLASH with higher access speed.

[rte_sws_3981] The memory section used to store measurement values in shall be the the memory sections associated with the `swAddressMethod` enclosed in the `SwDataDefProps` of a measurement definition.

Since it's measurement data obviously this must be in RAM.

[rte_sws_3982] The memory section used to store calibration parameters in shall be the the memory sections associated with the `swAddressMethod` enclosed in the `SwDataDefProps` of a calibration parameter definition.

4.2.6.3.5.8 Default parameter values

CalprmComponentPrototype or AUTOSAR SW-Cs have to provide one default value for each CalprmElementPrototype. The RTE has to apply this default value for a calibration parameters for all instances of the CalprmElementPrototype if not explicitly an additional calibration parameters is provided. This additional provided parameter value which can be instance specific overrules the default value.

Example:

A CalprmComponentType with calibration parameters for several window lifters contains these 2 CalprmElementPrototypes:

- LeftHandDrive
 - shared
 - Boolean
 - interpretation 0=left/1=right
 - default value = 0
- InitialWindowSpeed
 - instanceSpecific
 - UInt8
 - interpretation 0 = 0m/s, 255 = 0.255m/s
 - default value = 100

If the RTE generator doesn't get explicitly values for "LeftHandDrive" and "InitialWindowSpeed" than he will apply

1. value 0 for "LeftHandDrive"
2. value 100 for first instance of "InitialWindowSpeed"
3. value 100 for second instance of "InitialWindowSpeed"

If the RTE generator only gets additionally the explicit value 1 for "LeftHandDrive" than he will apply

1. value 1 for "LeftHandDrive"
2. value 100 for first instance of "InitialWindowSpeed"
3. value 100 for second instance of "InitialWindowSpeed"

If the RTE generator gets additionally the explicit value 1 for "LeftHandDrive" and the value 123 for the first instance of "InitialWindowSpeed" than he will apply

1. value 1 for "LeftHandDrive"
2. value 123 for first instance of "InitialWindowSpeed"
3. value 100 for second instance of "InitialWindowSpeed"

[rte_sws_3966] The RTE generator shall apply the default calibration parameter value for all instances of a CalprmElementPrototype for which no explicit alternative value is provided.

[rte_sws_3967] If an explicit calibration parameter value is specified for a CalprmElementPrototype the RTE generator shall apply this value according to the specified instance.

4.2.6.3.6 Calibration parameters in NVRAM

Calibration parameters can be located in NVRAM too. One use case for this is to have the possibility to modify calibration parameters via a diagnosis service without need for special calibration tool.

To allow NVRAM calibration parameters to be accessed, NVRAM with statically allocated RAM buffer in form of PIM memory for the calibration parameters has to be defined. Support of NVRAM with temporarily associated RAM buffer is not possible.

[rte_sws_3937] The RTE shall support calibration parameters configured to be allocated in NVRAM. Access to these parameters shall be supported via access of associated permanent PIM buffer.

[rte_sws_3938] The RTE generator shall reject configurations with calibration parameters defined in NVRAM when no associated NVRAM buffer in form of PIM has been defined statically.

4.2.6.3.7 Calibration parameters in Basic-SW

Calibration parameters can be defined in AUTOSAR SW as well as in Basic-SW (respectively in complex drivers). So data emulation must be possible for them all requiring special SW constructs for data emulation with SW support as described above.

But how can Basic-SW (respectively a complex drivers) force the allocation of their calibration parameters? In principle there are 2 possibilities. Which one to chose is not focus of this RTE specification.

1. RTE provides the calibration parameter access
Basic-SW implementing an AUTOSAR Interface can define a CalprmComponentType with the CalprmelementPrototypes it requires, define an AUTOSAR Interface in plus and define constraints about connections between them.
2. Calibration parameter access invisible for RTE
Since multiple instantiation is not allowed for Basic-SW it's possible for Basic-SW to define own methods how calibration parameters are allocated. In case data emulation with SW-support is used, the Basic-SW (respectively complex driver) which handles emulation details and data exchange with external calibration tools

then has to deal with two emulation methods at once: The one the RTE uses and the other one the Basic-SW practices.

4.3 Communication Models

AUTOSAR supports two basic communication patterns: Client-Server and Sender-Receiver. AUTOSAR software-components communicate through well defined ports and the behavior is statically defined by attributes. Some attributes are defined on the modeling level and others are closely related to the network topology and must be defined on the implementation level.

The RTE provides the implementation of these communication patterns. For inter-ECU communication the RTE uses the functionalities provided by COM. For intra-ECU communication the RTE can use the services of COM, but may as well implement the functionality on its own if that is more efficient.

With Sender-Receiver communication there are two main principles: Data Distribution and Event Distribution. When data is distributed, the last received value is of interest (last-is-best semantics). When events are distributed the whole history of received events is of interest, hence they must be queued on receiver side. Therefore an 'is-Queued' attribute of the data element is used to distinguish between Data and Event Distribution.² If a data element has event semantics, the isQueued attribute is set to true, if the data element has data semantics, the isQueued attribute is set to false.

[rte_sws_5508] The RTE generator shall reject the configuration when an r-port is connected to an r-port or a p-port is connected to a p-port with an AssemblyConnectorPrototypes or a ServiceConnectorPrototype.

For example, a required port (r-port) of a component typed by an AUTOSAR sender-receiver interface can read data elements of this interface. A provide port (p-port) of a component typed by an AUTOSAR sender-receiver interface can write data elements of this interface.

[rte_sws_7006] The RTE generator shall reject the configuration when an r-port is connected to a p-port or a p-port is connected to an r-port with a DelegationConnectorPrototypes.

4.3.1 Sender-Receiver

4.3.1.1 Introduction

Sender-receiver communication involves the transmission and reception of signals consisting of atomic data elements that are sent by one component and received by one or more components. A sender-receiver interface can contain multiple data elements.

²The isQueued attribute corresponds to the VFB attribute INFORMATION_TYPE.

Sender-receiver communication is one-way - any reply sent by the receiver is sent as a separate sender-receiver communication.

[rte_sws_5508] The RTE generator shall reject the configuration when an r-port is connected to an r-port or a p-port is connected to a p-port. A require port (r-port) of a component typed by an AUTOSAR sender-receiver interface can read data elements of this interface. A provide port (p-port) of a component typed by an AUTOSAR sender-receiver interface can write data elements of this interface.

4.3.1.2 Receive Modes

The RTE supports multiple receive modes for passing data to receivers. The four possible receive modes are:

- **“Implicit data read access”** – when the receiver’s runnable executes it shall have access to a “copy” of the data that remains unchanged during the execution of the runnable.

[rte_sws_6000] For data elements specified with implicit data read access, the RTE shall make the receive data available to the runnable through the semantics of a copy [RTE00128].

[rte_sws_6001] For data elements specified with implicit data read access the receive data shall not change during execution of the runnable [RTE00128].

When “implicit data read access” is used the RTE is required to make the data available as a “copy”. It is not necessarily required to use a unique copy for each runnable. Thus the RTE may use a unique copy of the data for each runnable entity or may, if several runnables (even from different components) need the same data, share the same copy between runnables. Runnable entities can only share a copy of the same data when the scheduling structure can make sure the contents of the data is protected from modification by any other party.

[rte_sws_6004] The RTE shall read the data elements specified with implicit data read access before the associated runnable entity is invoked [RTE00128].

Complex data types shall be handled in the same way as primitive data types, i.e. RTE shall make a “copy” available for the runnable.

[rte_sws_6003] The “implicit data read access” receive mode shall be valid for all categories of runnable entity (i.e. 1A, 1B and 2).[RTE00134].

- **“Explicit data read access”** – the RTE generator creates a non-blocking API call to enable a receiver to poll (and read) data. This receive mode is an “explicit” mode since an explicit API call is invoked by the receiver.

The explicit “data read access” receive mode is only valid for category 1B or 2 runnable entities [RTE00134].

- **“wake up of wait point”** – the RTE generator creates a blocking API call that the receiver invokes to read data.

[rte_sws_6002] The “wake up of wait point” receive mode shall support a timeout to prevent infinite blocking if no data is available [RTE00109].

The “wake up of wait point” receive mode is inherently only valid for a category 2 runnable entity.

A category 2 runnable entity is required since the implementation may need to suspend execution of the caller if no data is available.

- **“activation of runnable entity”** – the receiving runnable entity is invoked automatically by the RTE whenever new data is available. To access the new data, the runnable entity either has to use “implicit data read access” or “explicit data read access”, i.e. invoke an `Rte_Read` or `Rte_Receive` call, depending on the input configuration. This receive mode differs from “implicit data read access” since the receiver is invoked by the RTE in response to a `DataReceivedEvent`.

[rte_sws_6007] The “activation of runnable entity” receive mode shall be valid for category 1A, 1B and 2 runnable entities [RTE00134].

The validity of receive modes in conjunction with different categories of runnable entity is summarized in Table 4.4.

Receive Mode	Cat 1A	Cat 1B	Cat 2
Implicit Data Read Access	Yes	Yes	Yes
Explicit Data Read Access	No	Yes	Yes
Wake up of wait point	No	No	Yes
Activation of runnable entity	Yes	Yes	Yes

Table 4.4: Receive mode validity

The category of a runnable entity is not an inherent property but is instead determined by the features of the runnable. Thus the presence of explicit API calls makes the runnable at least category 1B and the presence of a wait point forces the runnable to be category 2.

4.3.1.2.1 Applicability

The different receive modes are not just used for receivers in sender-receiver communication. The same semantics are also applied in the following situations:

- **Success feedback** – The mechanism used to return transmission acknowledgments to a component. See Section 5.2.6.7.
- **Asynchronous client-server result** – The mechanism used to return the result of an asynchronous client-server call to a component. See Section 5.7.5.3.

4.3.1.2.2 Representation in the Software Component Template

The following list serves as a reference for how the RTE Generator determines the Receive Mode from its input [RTE00109]. Note that references to “the DataElement-Prototype” within this sub-section will implicitly mean “the DataElementPrototype for which the API is being generated”.

- “**wake up of wait point**” – A *DataReceivePoint* references a *DataElementPrototype* and a *WaitPoint* references a *DataReceivedEvent* which in turn references the same *DataElementPrototype*.
- “**activation of runnable entity**” – a *DataReceivedEvent* references the *DataElementPrototype* and a runnable entity to start when the data is received.
- “**explicit data read access**” – A *DataReceivePoint* references the *DataElementPrototype*.
- “**implicit data read access**” – A *DataReadAccess* references the *DataElementPrototype*.

It is possible to combine certain access methods; for example ‘activation of runnable entity’ can be combined with ‘explicit’ or ‘implicit’ data read access (indeed, one of these pairings is necessary to cause API generation to actually *read* the datum) but it is an input error if ‘activation of runnable entity’ and ‘wake up of wait point’ are combined (i.e. a *WaitPoint* references a *DataReceivedEvent* that references a runnable entity). It is also possible to specify both implicit and explicit data read access simultaneously.

For details of the semantics of “implicit data read access” and “explicit data read access” see Section 4.3.1.5.

4.3.1.3 Multiple Data Elements

A sender-receiver interface can contain one or more data elements. The transmission and reception of elements is independent – each data element, eg. AUTOSAR signal, can be considered to form a separate logical data channel between the “provide” port and a “require” port.

[rte_sws_6008] Each data element in a sender-receiver interface shall be sent separately [RTE00089].

Example 4.2

Consider an interface that has two data elements, `speed` and `freq` and that a component template defines a provide port that is typed by the interface. The RTE generator will then create two API calls; one to transmit `speed` and another to transmit `freq`.

Where it is important that multiple data elements are sent simultaneously they should be combined into a complex data structure (Section 4.3.1.11.1). The sender then cre-

ates an instance of the data structure which is filled with the required data before the RTE is invoked to transmit the data.

4.3.1.3.1 Initial Values

[rte_sws_6009] For each data element in an interface specified with data semantics (`isQueued = false`), the RTE shall support the `initValue` attribute [RTE00108].

The `initValue` attribute is used to ensure that AUTOSAR software-components always access valid data even if no value has yet been received. This information is required for both inter-ECU and intra-ECU communication. For inter-ECU communication initial values can be handled by COM but for intra-ECU communication RTE has to guarantee that `initValue` is handled.

The specification of an init value is mandatory for each data element prototype with `isQueued = FALSE`, see [17].

[rte_sws_6010] When `isQueued` is specified as false, the RTE shall use any specified initial value to prevent the receiver performing calculations based on invalid (i.e. uninitialized) values [RTE00107].

The above requirement ensures that RTE API calls return the initialized value until a “real” value has been received, possibly via the communication service. The requirement does *not* apply when the `isQueued` attribute is set to true, i.e. when “event” semantics are used since the implied state change when the event data is received will mean that the receiver will not start to process invalid data and would therefore never see the initialized value.

[rte_sws_4500] An initial value cannot be specified when the `isQueued` attribute is specified as true [RTE00107].

For senders, an initial value is not used directly by the RTE (since an AUTOSAR SW-C must supply a value using `Rte_Send`) however it may be needed to configure the communication service - for example, an un-initialised signal can be transmitted if multiple signals are mapped to a single frame and the communication service transmits the whole frame when any contained signal is sent by the application. Note that it is not the responsibility of the RTE generator to configure the communication service.

It is permitted for an initial value to be specified for either the sender or receiver. In this case the same value is used for both sides of the communication.

[rte_sws_4501] If in context of one ECU a sender specifies an initial value and the receiver does not (or *vice versa*) the same initial value is used for both sides of the communication [RTE00108].

It is also permitted for both sender and receiver to specify an initial value. In this case it is defined that the receiver’s initial value is used by the RTE generator for both sides of the communication.

[rte_sws_4502] If in context of one ECU both receiver and sender specify an initial value the specification for the *receiver* takes priority [RTE00108].

4.3.1.4 Multiple Receivers and Senders

Sender-receiver communication is not restricted to communication connections between a single sender and a single receiver. Instead, sender receiver communication connection can have multiple senders ('n:1' communication) or multiple receivers ('1:m' communication) with the restrictions that multiple senders are not allowed for `mode switch notifications`, see `metamodel restriction rte_sws_2670`.

The RTE does not impose any co-ordination on senders – the behavior of senders is independent of the behavior of other senders. For example, consider two senders *A* and *B* that both transmit data to the same receiver (i.e. 'n:1' communication). Transmissions by either sender can be made at any time and there is no requirement that the senders co-ordinate their transmission. However, while the RTE does not impose any co-ordination on the senders it does ensure that simultaneous transmissions do not conflict.

In the same way that the RTE does not impose any co-ordination on senders there is no co-ordination imposed on receivers. For example, consider two receivers *P* and *Q* that both receive the same data transmitted by a single sender (i.e. '1:m' communication). The RTE does not guarantee that multiple receivers see the data simultaneously even when all receivers are on the same ECU.

4.3.1.5 Implicit and Explicit Data Reception and Transmission

[rte_sws_6011] The RTE shall support 'explicit' and 'implicit' data reception and transmission.

Implicit data access transmission means that a runnable does not actively initiate the reception or transmission of data. Instead, the required data is received automatically when the runnable starts and is made available for other runnables at the earliest when it terminates.

Explicit data reception and transmission means that a runnable employs an explicit API call to send or receive certain data elements. Depending on the category of the runnable and on the configuration of the according ports, these API calls can be either blocking or non-blocking.

4.3.1.5.1 Implicit

DataReadAccess

For the implicit reading of data, called *DataReadAccess* [RTE00128], the data is made available when the runnable starts using the semantics of a copy operation and the RTE ensures that the 'copy' will not be modified until after the runnable terminates.

When a runnable *R* is started, the RTE reads all data elements marked with 'DataReadAccess', if the data elements may be changed by other runnables a copy is created that will be available to runnable *R*. The runnable *R* can read the data element by using the RTE APIs for implicit read (see the API description in Sect. 5.6.15). That way, the data is guaranteed not to change (e.g. by write operations of other runnables) during the entire lifetime of *R*. If several runnables (even from different components) need the data, they can share the *same* buffer. This is only applicable when the scheduling structure can make sure the contents of the data is protected from modification by any other party.

Note that this concept implies that the runnable does in fact terminate. Therefore, while *DataReadAccess* is allowed for category 1A and 1B runnable entities as well as category 2 only the former are guaranteed to have a finite execution time. A category 2 runnables that runs forever will not see any updated data.

DataReadAccess is only allowed for DataElement-Prototypes with their `isQueued` attribute set to false (`rte_sws_3012`).

DataWriteAccess

Implicit sending, called *DataWriteAccess* [RTE00129], is the opposite concept. Data elements marked as 'DataWriteAccess' are sent by the RTE after the runnable terminates. The runnable can write the data element by using the RTE APIs for implicit write (see the API description in Sect. 5.6.16 and 5.6.17). The sending is independent from the position in the execution flow in which the `Rte_IWrite` is performed inside the Runnable. When performing several write accesses during runnable execution to the same data element, only the last one will be recognized. Here we have a last-is-best semantics.

Note:

If *DataWriteAccess* is specified for a certain data element, but no RTE API for implicit write of this data element is called during an execution of the runnable, an undefined value is written back when the runnable terminates.

[rte_sws_3570] For *DataWriteAccess* the RTE shall make the send data available to others (other runnables, other AUTOSAR SWCs, Basic SW, ..) with the semantics of a copy [RTE00129].

[rte_sws_3571] For *DataWriteAccess* the RTE shall make the send data available to others (other runnables, other AUTOSAR SWCs, Basic SW, ..) at the earliest when the runnable returns (exits the 'Running' state) [RTE00129].

[rte_sws_3572] For *DataWriteAccesses* several accesses to the same data element performed inside a runnable during one runnable execution shall lead to only one transmission of the data element [RTE00129].

[rte_sws_3573] If several DataWriteAccesses to the same data element are performed inside a runnable during the runnable execution, the RTE shall use the last value written. (last-is-best semantics) [RTE00129]

DataWriteAccess is only sensible for runnable entities that are guaranteed to terminate, i.e. category 1A and 1B. It is use DataWriteAccess for a category 2 runnable but if they do not terminate then no data write-back will occur.

[rte_sws_3574] DataWriteAccess shall be valid for all categories of runnable entity [RTE00134].

To get common behavior in RTEs from different suppliers further requirements defining the semantic of implicit communication exist:

[rte_sws_3954] Buffers used by the RTE to contain data copies for implicit communication shall have task wide scope.

Requirement rte_sws_3954 means that all runnable entities mapped to a task that access a data element using DataReadAccess and/or DataWriteAccess access the same buffers.

[rte_sws_3598] For implicit communication, a single shared read/write buffer shall be used when no runnable entity mapped to the task has both DataReadAccess and DataWriteAccess to the same data element.

[rte_sws_3955] For implicit communication, separate read and write buffers shall be used when at least one runnable entity mapped to the task has both DataReadAccess and DataWriteAccess to the same data element.

[rte_sws_3599] For implicit communication all readers (runnables that perform DataReadAccess) within a task shall access the same buffer.

[rte_sws_3953] For implicit communication all writers (runnables that perform DataWriteAccess) within a task shall access the same buffer.

The content of a shared buffer (see rte_sws_3598) is not guaranteed to stay constant during the whole task since a writer will change the shared copy and hence readers mapped in the task after the writer will access the updated copy. When buffers are shared, written data is visible to other runnables within the same execution of the task. However since no runnable within the task will both read and write the buffer (rte_sws_3598) consistency *within a runnable* is ensured.

When separate buffers used for implicit communication (see rte_sws_3955) any data written by a runnable is not visible (to either other runnables or to the writing runnable) until the data is written back after the runnable has terminated. For runnables within the same task it will not be visible until the next task execution.

[rte_sws_3956] The content of a task specific buffer used for DataReadAccess shall be filled with actual data by a copy action at the begin of the task.

[rte_sws_3957] The content of a buffer modified by `DataWriteAccess` in one task shall be made available to runnable entities using `DataReadAccess` allocated in other tasks after the execution of the last runnable mapped to the task.

Note:

It's the semantic of implicit communication that a `DataWriteAccess` is interpreted as writing the whole data element.

4.3.1.5.2 Explicit

The behavior of explicit reception depends on the category of the runnable and on the configuration of the according ports.

An explicit API call can be either non-blocking or blocking. If the call is non-blocking (i.e. there is a `DataReceivePoint` referencing the `DataElementPrototype` for which the API is being generated, but no `WaitPoint` referencing the `DataReceivePoint`), the API call immediately returns the next value to be read and, if the communication is queued (event reception), it removes the data from the receiver-side queue, see Section 4.3.1.10

[rte_sws_6012] A non-blocking RTE API “read” call shall indicate if no data is available [RTE00109].

In contrast, a blocking call (i.e. there is a `WaitPoint` referencing the `DataReceivePoint` for which the API is being generated) will suspend execution of the caller until new data arrives (or a timeout occurs) at the according port. When new data is received, the RTE resumes the execution of the waiting runnable. ([RTE00092])

To prevent infinite waiting, a blocking RTE API call can have a timeout applied. The RTE monitors the timeout and if it expires without data being received returns a particular error status.

[rte_sws_6013] A blocking RTE API “read” call shall indicate the expiry of a timeout [RTE00069].

The “timeout expired” indication also indicates that no data was received before the timeout expired.

Blocking reception of data (“wake up of wait point” receive mode as described in Section 4.3.1.2) is only applicable for category 2 runnables whereas non-blocking reception (“explicit data read access” receive mode) can be employed by runnables of category 2 or 1B. Neither blocking nor non-blocking explicit reception is applicable for category 1A runnable because they must not invoke functions with unknown execution time (see table 4.4).

[rte_sws_6016] The RTE API call for explicit sending (`DataSendPoint`, [RTE00098]) shall be non-blocking.

Using this API call, the runnable can explicitly send new values of the according data element.

Explicit writing is valid for runnables of category 1b and 2 only. Explicit writing is not allowed for a category 1A runnable since these require API calls with constant execution time (i.e. macros).

Although the API call for explicit sending is non-blocking, it is possible for a category 2 runnable to block waiting for a notification whether the (explicit) send operation was successful. This is specified by the `AcknowledgementRequest` attribute and occurs by a separate API call `Rte.Feedback`. If the feedback method is `'wake_up_of_wait_point'`, the runnable will block and be resumed by the RTE either when a positive or negative acknowledgement arrives or when the timeout associated with the wait point expires.

4.3.1.5.3 Concepts of data access

Tables 4.5 and 4.6 summarize the characteristics of implicit versus explicit data reception and transmission.

Implicit Read	Explicit Read
Receiving of data element values is performed only once when runnable starts	Runnable decides when and how often a data element value is received
Values of data elements do not change while runnable is running.	Runnable can always decide to receive the latest value
Several API calls to the same signal always yield the same data element value	Several API calls to the same signal may yield different data element values
Runnable must terminate (all categories)	Runnable is of cat. 1B or 2

Table 4.5: Implicit vs. explicit read

Implicit Write	Explicit Write
Sending of data element values is only done once after runnable returns	Runnable can decide when sending of data element values is done via the API call
Several usages of the API call inside the runnable cause only one data element transmission	Several usages of the API call inside the runnable cause several transmissions of the data element content. (Depending on the behavior of COM, the number of API calls and the number of transmissions are not necessarily equal.)
Runnable must terminate (all categories)	Runnable is cat. 1B or 2

Table 4.6: Implicit vs. explicit write

4.3.1.6 Transmission Acknowledgement

When `AcknowledgementRequest` is specified, the RTE will inform the sending component if the signal has been sent correctly or not. Note that there is no insurance that the signal has actually been *received* correctly by the corresponding receiver AUTOSAR software-component. Thus, only the RTE on the sender side is involved in supporting `AcknowledgementRequest`.

In case of mode switch communication (see Section 4.4), the communication is local to one ECU. The transmission acknowledgement will be sent, when the mode switch is executed by the RTE, see `rte.sws_2587`.³

³Currently, no mode switch acknowledgement is defined. If a mode switch acknowledgement will be defined in future releases, it shall be used instead of the transmission acknowledgement.

[rte_sws_5504] The RTE shall support the use of `AcknowledgementRequest` independently for each data item of an AUTOSAR software-component's AUTOSAR interface [RTE00122].

[rte_sws_5506] The RTE generator shall reject specification of the `AcknowledgementRequest` attribute for transmission acknowledgement for 1:n communication [RTE00125], except for mode switch communication. Restriction: In some cases, when more than one receiver is connected via one physical bus, this can not be discovered by the RTE generator.

The result of the feedback can be collected using “wake up of wait point”, “explicit data read access” or “activation of runnable entity”.

The `AcknowledgementRequest` attribute allows to specify a timeout.

[rte_sws_3754] If `AcknowledgementRequest` is specified, the RTE shall ensure that timeout monitoring is performed, regardless of the receive mode of the acknowledgement.

For inter-ECU communication, AUTOSAR COM provides the necessary functionality, for intra-ECU communication, the RTE has to implement the timeout monitoring.

If a `WaitPoint` is specified to collect the acknowledgement, two timeout values have to be specified, one for the `AcknowledgementRequest` and one for the `WaitPoint`.

[rte_sws_3755] If different timeout values are specified for the `AcknowledgementRequest` for a `DataElementPrototype` and for the `WaitPoint` associated with the `DataSendCompletedEvent` for the `DataSendPoint` for that `DataElementPrototype`, the configuration shall be rejected by the RTE generator.

The `DataSendCompletedEvent` associated with the `DataSendPoint` for a `DataElementPrototype` shall indicate that the transmission was successful or that the transmission was not successful. The status information about the success of the transmission shall be available as the return value of the generated RTE API call.

[rte_sws_3756] For each transmission of a `DataElementPrototype` only one acknowledgement shall be passed to the sending component by the RTE. The acknowledgement indicates either that the transmission was successful or that the transmission was not successful.

[rte_sws_3757] The status information about the success or failure of the transmission shall be available as the return value of the RTE API call to retrieve the acknowledgement.

[rte_sws_3604] The status information about the success or failure of the transmission shall be buffered with last-is-best semantics. When a data item is sent, the status information is reset.

`rte_sws_3604` implies that once the `DataSendCompletedEvent` has occurred, repeated API calls to retrieve the acknowledgement shall always return the same result until the next data item is sent.

[rte_sws_3758] If the timeout value of the AcknowledgementRequest is 0, no timeout monitoring shall be performed.

4.3.1.7 Communication Time-out

When sender-receiver communication is performed using some physical network there is a chance this communication may fail and the receiver does not get an update of data (in time or at all). To allow the receiver of a `data element` to react appropriately to such a condition the SW-C template allows the specification of a time-out which the infrastructure shall monitor and indicate to the interested software components.

A “data element” is the actual information exchanged in case of sender-receiver communication. In the COM specification this is represented by a `ComSignal`. In the SW-C template a data element is represented by the instance of a `DataElementPrototype`.

[rte_sws_5020] When present, the `aliveTimeout` attribute⁴ `rte_sws_in_0067` enables the monitoring of the timely reception of the `data element` with data semantics (`is-Queued = false`) transmitted over the network.

The monitoring functionality is provided by the COM module, the RTE transports the event of reception time-outs to software components as “data element outdated”. The software components can either subscribe to that event (activation of runnable entity) or get that situation passed by the implicit and explicit status information (using API calls).

[rte_sws_5021] If `aliveTimeout` is present, but the communication is local to the ECU, time-out monitoring is disabled and no notification of the software components will occur.

Therefore the Software Component shall not rely in its functionality on the time-out notification, because for local communication the notification will never occur. Time-out notification is intended as pure error reporting.

[rte_sws_3759] If the `aliveTimeout` attribute is 0, no timeout monitoring shall be performed.

[rte_sws_5022] If a time-out has been detected, the last correctly received value shall be provided to the software components (preserving the last-is-best-semantics, see Section 4.3.1.10.1).

The time-out support (called “deadline monitoring” in COM) provided by COM has some restrictions which have to be respected when using this mechanism. Since the COM module is configured based on the System Description the restrictions mainly arise from the `data element` to I-PDU mapping. This already has to be considered when developing the System Description and the RTE Generator can only provide warnings when inconsistencies are detected. Therefore the RTE Generator needs to have access to the configuration information of COM.

⁴This attribute is called “LIVELIHOOD” in the VFB specification

In case time-out is enabled on a `data element` with update bit, there shall be a separate time-out monitoring for each `data element` with an update bit [COM292].

There shall be an I-PDU based time-out for `data elements` without an update bit [COM290]. For all data elements without update bits within the same I-PDU, the smallest configured time-out of the associated data elements is chosen as time-out for the I-PDU[COM291]. The notification from COM to RTE is performed per data element.

In case one `data element` coming from COM needs to be distributed to several SW-Components the SW-C template allows to specify different `aliveTimeout` values at each Port. But COM does only support one `aliveTimeout` value per `data element`, therefore the smallest `aliveTimeout` value shall be used for the notification of the time-out to several SW-Components.

4.3.1.8 Data Element Invalidation

The Software Component template allows to specify whether a `data element`, defined in an AUTOSAR Interface, can be invalidated by the sender. The communication infrastructure shall provide means to set a data element to invalid and also indicate an invalid data element to the receiving software components. This functionality is called “data element invalidation”.

[rte_sws_5024] On sender side the `canInvalidate` attribute `rte_sws_in_5023` (when present) enables the invalidation support for this `data element`. The actual value used to represent the invalid data element shall be specified in the Data Semantics part of the data element definition defined in `rte_sws_in_5031`⁵.

[rte_sws_5032] On receiver side the `handleInvalid` attribute `rte_sws_in_5050` of the `UnqueuedReceiverComSpec` specifies how to handle the reception of the invalid value.

[rte_sws_5033] Data element invalidation is only supported for data elements with the `isQueued` attribute set to `false` `rte_sws_in_45`.

The API to set a `data element` to invalid shall be provided to the runnable entities on data element level.

In case an invalidated data element is received a software component can be notified using the activation of runnable entity. If an invalidated data element is read by the SW-C the invalid status shall be indicated in the status code of the API.

4.3.1.8.1 Data Element Invalidation in case of Inter-ECU communication

Sender:

If `canInvalidate` is enabled and the communication is Inter-ECU:

⁵When `canInvalidate` is enabled but there is no `invalidValue` specified it is considered an invalid configuration.

- explicit data transmission: data element invalidation will be performed by COM (COM needs to be configured properly).
- implicit data transmission: data element invalidation will be performed by RTE.

Receiver:

[rte_sws_5026] If a data element has been received invalidated in case of Inter-ECU communication and the attribute `handleInvalid` `rte_sws_in_5050` is set to `keep` – the query of the value shall return the value provided by COM together with an indication of the invalid case. Then the reception of the invalid value will be handled as an error and the activation of runnable entities can be performed using the `DataReceiveErrorEvent`.

[rte_sws_5048] If a data element has been received invalidated in case of Inter-ECU communication and the attribute `handleInvalid` `rte_sws_in_5050` is set to `replace` – COM shall be configured to perform the “invalid value substitution” (`Com_DataInvalidAction` is `Replace` [COM314]) with the `initValue`. Then the reception will be handled as if a valid value would have been received (activation of runnable entities using the `DataReceivedEvent`).

4.3.1.8.2 Data Element Invalidation in case of Intra-ECU communication

Sender:

[rte_sws_5025] If `canInvalidate` is enabled, and the communication is Intra-ECU, data element invalidation can be implemented by the RTE or the RTE may utilize the implementation of the AUTOSAR COM module.

In case of implicit data transmission the RTE shall always implement the data element invalidation and therefore provide an API to set the data element’s value to the invalid value. The actual invalid value is specified in the SW-C template `rte_sws_in_5031`.

Receiver:

[rte_sws_5030] If a data element has been invalidated in case of Intra-ECU communication and the attribute `handleInvalid` `rte_sws_in_5050` is set to `keep` – the query of the value shall return the same value as if COM would have handled the invalidation (copy COM behavior). Then the reception of the invalid value will be handled as an error and the activation of runnable entities can be performed using the `DataReceiveErrorEvent`.

[rte_sws_5049] If a data element has been received invalidated in case of Intra-ECU communication and the attribute `handleInvalid` `rte_sws_in_5050` is set to `replace` – RTE shall perform the “invalid value substitution” with the `initValue`. Then the reception will be handled as if a valid value would have been received (activation of runnable entities using the `DataReceivedEvent`).

4.3.1.9 Filters

By means of the `filter` attribute [RTE00121] an additional filter layer can be added on the receiver side. *Value-based* filters can be defined, i.e. only signal values fulfilling certain conditions are made available for the receiving component. The possible filter algorithms are taken from OSEK COM version 3.0.2. They are listed in the meta model (see [17], Sect. 'Communication specification of data filters'). According to the SW-C template [17], filters are only allowed for signals that are compatible to C language unsigned integer types (i.e. characters, unsigned integers and enumerations). Thus, filters cannot be applied to complex data types like records or arrays.

[rte_sws_5503] The RTE shall provide value-based filters on the receiver-side as specified in the SW-C template [17], Section 'Communication specification of data filters'.

[rte_sws_5500] For inter-ECU communication, the RTE shall use the filter implementation of the COM layer [RTE00121]. For intra-ECU communication, the RTE can use the filter implementation of COM, but may also implement the filters itself for efficiency reasons, without using COM.

[rte_sws_5501] The RTE shall support a different filter specification for each data element in a component's AUTOSAR interface [RTE00121].

4.3.1.10 Buffering

[rte_sws_2515] The buffering of sender-receiver communication shall be done on the receiver side. This does not imply that COM does no buffering on the sender side. On the receiver side, two different approaches are taken for the buffering of 'data' and of 'events', depending on the value of the `isQueued` attribute of the data element.

4.3.1.10.1 Last-is-Best-Semantics for 'data' Reception

[rte_sws_2516] On the receiver side, the buffering of 'data' (`isQueued = false`) shall be realized by the RTE by a single data set for each data element instance.

The use of a single data set provides the required semantics of a single element queue with overwrite semantics (new data replaces old). Since the RTE is required to ensure data consistency, the generated RTE should ensure that non-atomic reads and writes of the data set (e.g. for complex data) are protected from conflicting concurrent access. RTE may use lower layers like COM to implement the buffer.

[rte_sws_2517] Depending on the ports attributes, the RTE shall initialize this data set with a startup value.

[rte_sws_2518] Implicit or explicit read access shall always return the last received data.

Requirement `rte_sws_2518` applies whether or not there is a `DataReceivedEvent` referencing the `DataElementPrototype` for which the API is being generated.

[rte_sws_2519] Explicit read access shall be non blocking in the sense that it does not wait for new data to arrive. The RTE shall provide mutual exclusion of read and write accesses to this data, e.g., by `ExclusiveAreas`.

[rte_sws_2520] When new data is received, the RTE shall silently discard the previous value of the data, regardless of whether it was read or not.

4.3.1.10.2 Queueing for 'event' Reception

The application of event semantics implies a state change. Events usually have to be handled. In many cases, a loss of events can not be tolerated. Hence the `isQueued` attribute is set to true to indicate that the received 'events' have to be buffered in a queue.

[rte_sws_2521] The RTE shall implement a receive queue for each event-like data element (`isQueued = true`) of a receive port.

The `queueLength` attribute of the `EventReceiverComSpec` referencing the event assigns a constant length to the receive queue.

[rte_sws_2522] The events shall be written to the end of the queue and read (consuming) from the front of the queue (i.e. the queue is first-in-first-out).

[rte_sws_2523] If a new event is received when the queue is already filled, the RTE shall discard the received event and set an error flag.

[rte_sws_2524] The error flag shall be reset during the next explicit read access on the queue. In this case, the status value `RTE_E_LOST_DATA` shall be presented to the application together with the data.

[rte_sws_2525] If an empty queue is polled, the RTE shall return with a status `RTE_E_NO_DATA` to the polling function, (see chap. 5.5.1).

The minimum size of the queue is 1.

[rte_sws_2526] The RTE generator shall reject a `queueLength` attribute of an `EventReceiverComSpec` with a queue length ≤ 0 .

4.3.1.11 Operation

4.3.1.11.1 Inter-ECU Mapping

This section describes the mapping from `DataElementPrototypes` to COM signals or COM signal groups for sender-receiver communication. The mapping is described

in the input of the RTE generator, in the DataMapping section of the System Template [14].

If a DataElementPrototype is mapped to a COM signal or COM signal group but the communication is local, the RTE generator can use the COM signal/COM signal group for the transmission or it can use its own direct implementation of the communication for the transmission.

4.3.1.11.1.1 Primitive Data Types

[rte_sws_4504] If a data element is a primitive type and the communication is inter-ECU, the DataMappings element shall contain a mapping of the data element to at least one COM signal, else the missing data mapping shall be interpreted as an unconnected port.

The mapping defines all aspects of the signal necessary to configure the communication service, for example, the network signal endianness and the communication bus. The RTE generator only requires the COM signal handle id since this is necessary for invoking the COM API.

[rte_sws_4505] The RTE shall use the handle id of the corresponding COM signal when invoking the COM API for signals.

The actual COM handle id has to be gathered from the ECU configuration of the COM module. The input information `rte_sws_in_5079` is used to establish the link between the ComSignal of the COM module's configuration and the corresponding SignalInstance of the System Template.

4.3.1.11.1.2 Complex Data Types

When a data element is a complex type the RTE is required to perform more complex actions to marshal the data [RTE00091] than is the case for primitive data types.

The DataMappings element of the ECU configuration contains (or reference) sufficient information to allow the data item or operation parameters to be transmitted. The mapping indicates the COM signals or signal groups to be used when transmitting a given data item of a given port of a given software component instance within the composition.

[rte_sws_4506] If a data element is a complex type and the communication is inter-ECU, the DataMappings element shall contain a mapping of the data element to COM signals such that each element of the complex data type that is a primitive data type is mapped to a separate COM signal(s), else the missing data mapping shall be interpreted as an unconnected port.

[rte_sws_4507] If a data element is a complex type and the communication is inter-ECU, the DataMappings element shall contain a mapping of the data element to COM

signals such that each element of the complex data type that is itself a complex data type shall be recursively mapped to a primitive type and hence to a separate COM signal(s).

The above requirements have two key features; firstly, COM is responsible for endianness conversion (if any is required) of primitive types and, secondly, differing structure member alignment between sender and receiver is irrelevant since the COM signals are packed into I-PDUs by the COM configuration.

The DataMappings shall contain sufficient COM signals to map each primitive element⁶ of the AUTOSAR signal.

[rte_sws_4508] If a data element is a complex type and the communication is inter-ECU, the DataMappings element shall contain at least one COM signal for each primitive element of the AUTOSAR signal.

[rte_sws_2557]

1. Each signal that is mapped to an element of the same composite data item shall be mapped to the same signal group.
2. If two signals are not mapped to an element of the same composite data item, they shall not be mapped to the same signal group.
3. If a signal is not mapped to an element of a composite data item, it shall not be mapped to a signal group.

[rte_sws_5081] The RTE shall use the handle id of the corresponding COM signal group when invoking the COM API for signal groups.

The actual COM handle id has to be gathered from the ECU configuration of the COM module. The input information `rte_sws.in_5080` is used to establish the link between the `ComSignalGroup` of the COM module's configuration and the corresponding `SignalInstance` of the System Template.

4.3.1.11.2 Atomicity

[rte_sws_4527] The RTE is required to treat AUTOSAR signals transmitted using sender-receiver communication atomically [RTE00073]. To achieve this the "signal group" mechanisms provided by COM shall be utilized. See `rte_sws_2557` for the mapping.

The RTE decomposes the complex data type into single signals as described above and passes them to the COM module by using the COM API call `ComUpdateShadowSignal`. As this set of single signals has to be treated as atomic, it is placed in a "signal group". A signal group has to be placed always in a single I-PDU.

⁶An AUTOSAR signal that is a primitive data type contains exactly one primitive element whereas a signal that is a complex type contains one or more primitive elements.

Thus, atomicity is established. When all signals have been updated, the RTE initiates transmission of the signal group by using the COM API call `Com_SendSignalGroup`.

As would be expected, the receiver side is the exact reverse of the transmission side: the RTE must first call `Com_ReceiveSignalGroup` precisely once for the signal group and then call `Com_ReceiveShadowSignal` to extract the value of each signal within the signal group.

A signal group has the additional property that COM guarantees to inform the receiver by invoking a call-back about its arrival only after all signals belonging to the signal group have been unpacked into a shadow buffer.

4.3.1.11.3 Fan-out

Fan-out can be divided into two scenarios; “PDU fanout” where the same I-PDU is sent to multiple destinations and “signal fan-out” where the same signal, i.e. data element is sent in different I-PDUs to multiple receivers.

For Inter-ECU communication, the RTE does not perform PDU fan-out. Instead, the RTE invokes `Com_SendSignal` once for a primitive data element and expects the fan-out to multiple destinations to occur lower down in the AUTOSAR communication stack. However, it is necessary for the RTE to support “signal fan-out” since this cannot be performed by any lower level layer of the AUTOSAR communication stack.

[rte_sws_6023] For inter-ECU transmission of a primitive data type, the RTE shall invoke `Com_SendSignal` for each COM signal to which the primitive data element is mapped.

If the data element is a complex data type, RTE invokes `Com_UpdateShadowSignal` for each primitive element in the complex data type and each COM signal to which that primitive element is mapped, and `Com_SendSignalGroup` for each COM signal group to which the data element is mapped.

[rte_sws_4526] For inter-ECU transmission of complex data type, the RTE shall invoke `Com_UpdateShadowSignal` for each COM signal to which an element in the complex data type is mapped and `Com_SendSignalGroup` for each COM signal group to which the complex data element is mapped.

For intra-ECU transmission of data elements, the situation is slightly different; the RTE handles the communication (the lower layers of the AUTOSAR communication stack are not used) and therefore must ensure that the data elements are routed to all receivers.

[rte_sws_6024] For intra-ECU transmission of data elements, the RTE shall perform the fan-out to each receiver [RTE00028].

4.3.1.11.4 Fan-in

When receiving data from multiple senders in inter-ECU communication, either the RTE on the receiver side has to collect data received in different COM signals or COM signal groups and pass it to one receiver or the RTE on the sender side has to provide shared access to a COM signal or COM signal group to multiple senders. The receiver RTE, which has to handle multiple COM signals or signal groups, is notified about incoming data for each COM signal or COM signal group separately but has to ensure data consistency when passing the data to the receiver. The sender RTE, which has to handle multiple senders sharing COM signals or signal groups, has to ensure consistent access to the COM API, since COM API calls for the same signal are not reentrant.

[rte_sws_3760] If multiple senders use different COM signals or signal groups for inter-ECU transmission of a data element prototype with `isQueued = false` to a receiver, the RTE on the receiver side has to pass the last received value to the receiver component while ensuring data consistency.

[rte_sws_3761] If multiple senders use different COM signals or signal groups for inter-ECU transmission of a data element prototype with `isQueued = true` to a receiver, the RTE on the receiver side has to queue all incoming values while ensuring data consistency.

[rte_sws_3762] If multiple senders share COM signals or signal groups for inter-ECU transmission of a data element prototype to a receiver, the RTE on the sender side has to ensure that the COM API for those signals is not invoked concurrently.

For intra-ECU transmission, the RTE must handle the fan-in, which is already stated in `rte_sws_6024`.

4.3.1.11.5 Sequence diagrams of Sender Receiver communication

Figure 4.24 shows a sequence diagram of how Sender Receiver communication for data transmission and non-blocking reception may be implemented by RTE. The sequence diagram also shows the `Rte_Read` API behavior if an `initValue` is specified.

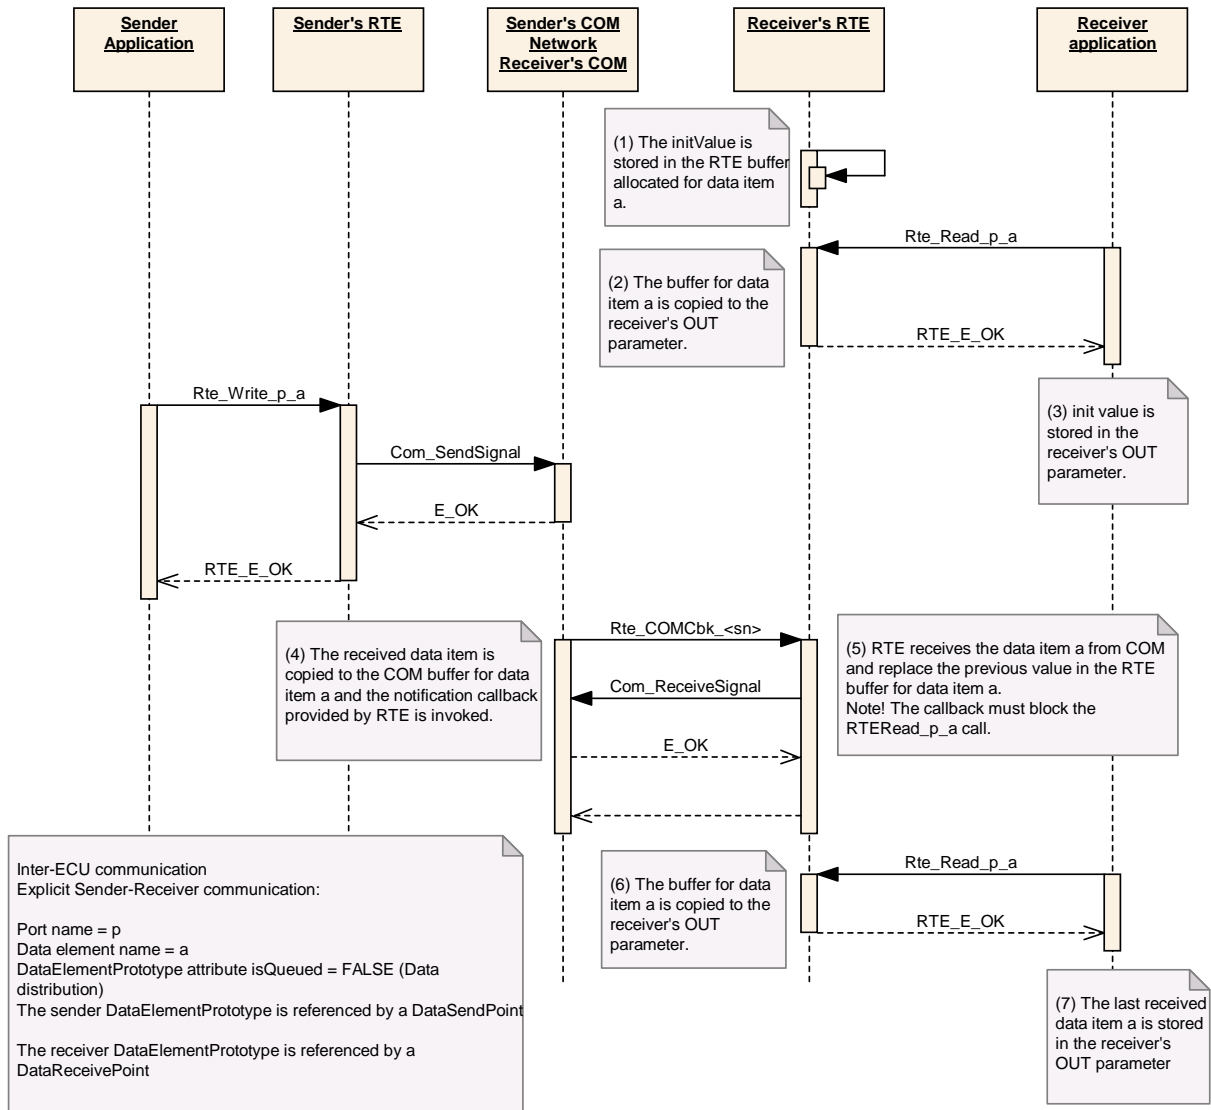


Figure 4.24: Sender Receiver communication with isQueued false and DataReceivePoint as reception mechanism

Figure 4.25 shows a sequence diagram of how Sender Receiver communication for event transmission and non-blocking reception may be implemented by RTE. The sequence diagram shows the `Rte_Receive` API behavior when the queue is empty.

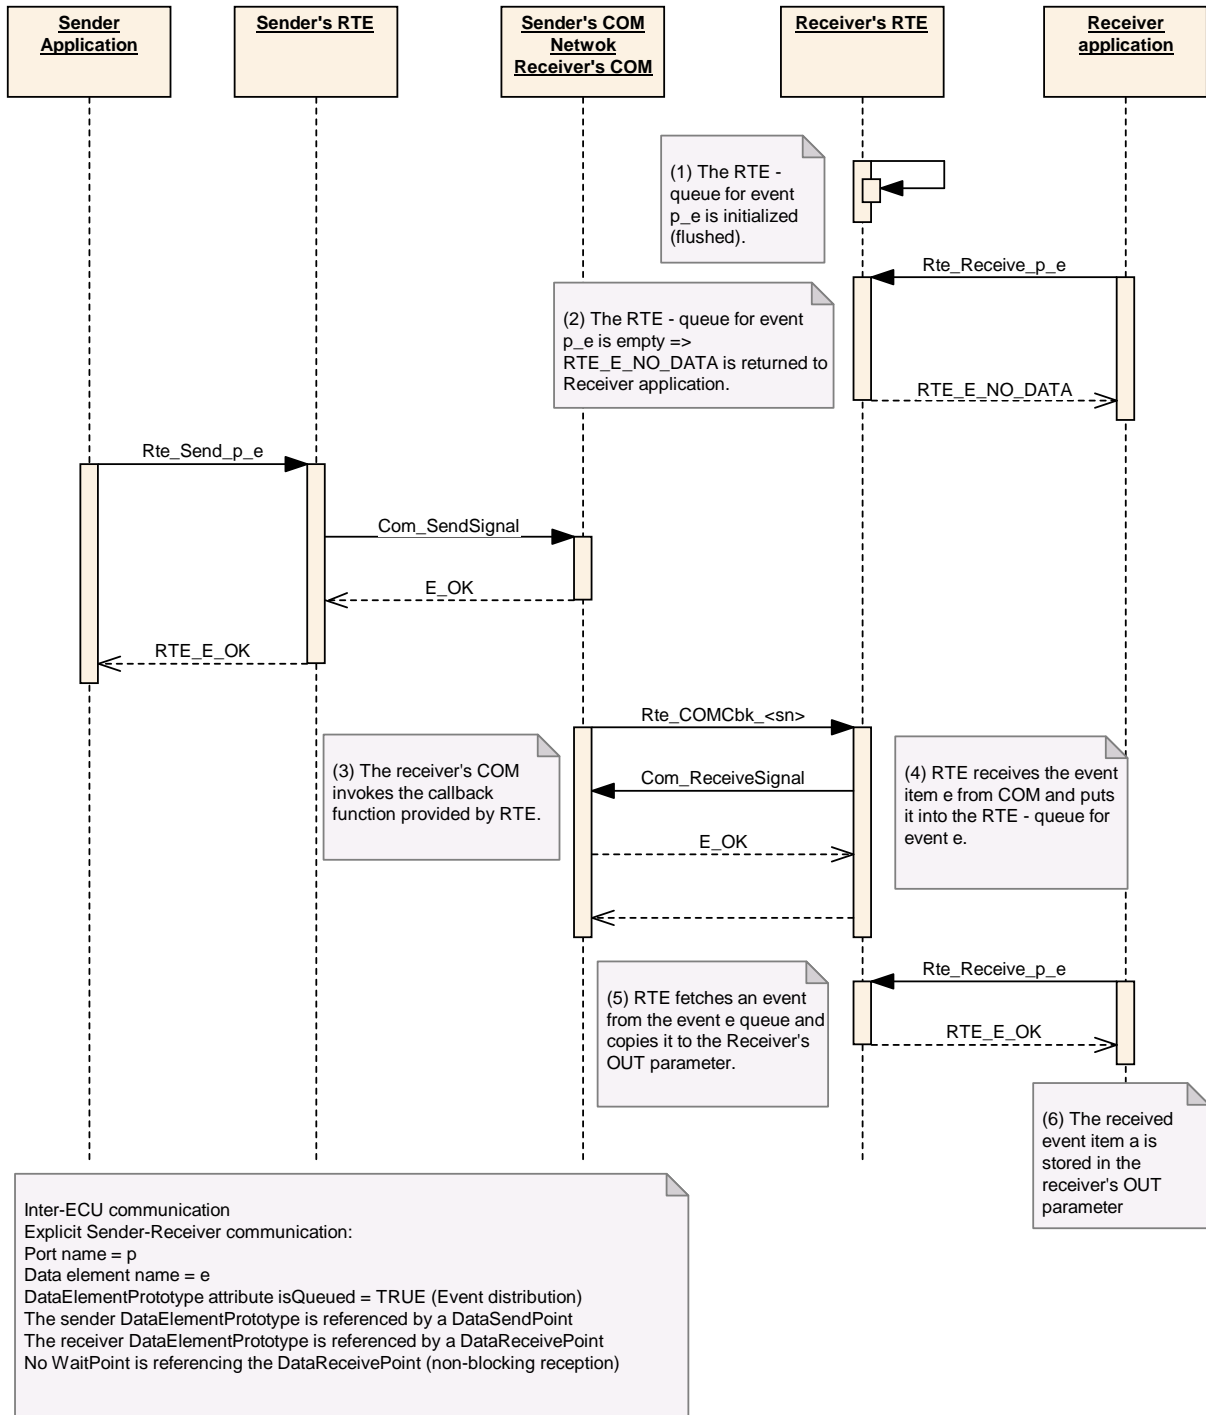


Figure 4.25: Sender Receiver communication with `isQueued` true and `DataReceivePoint` as reception mechanism

Figure 4.26 shows a sequence diagram of how Sender Receiver communication for event transmission and activation of runnable entity on the receiver side may be implemented by RTE.

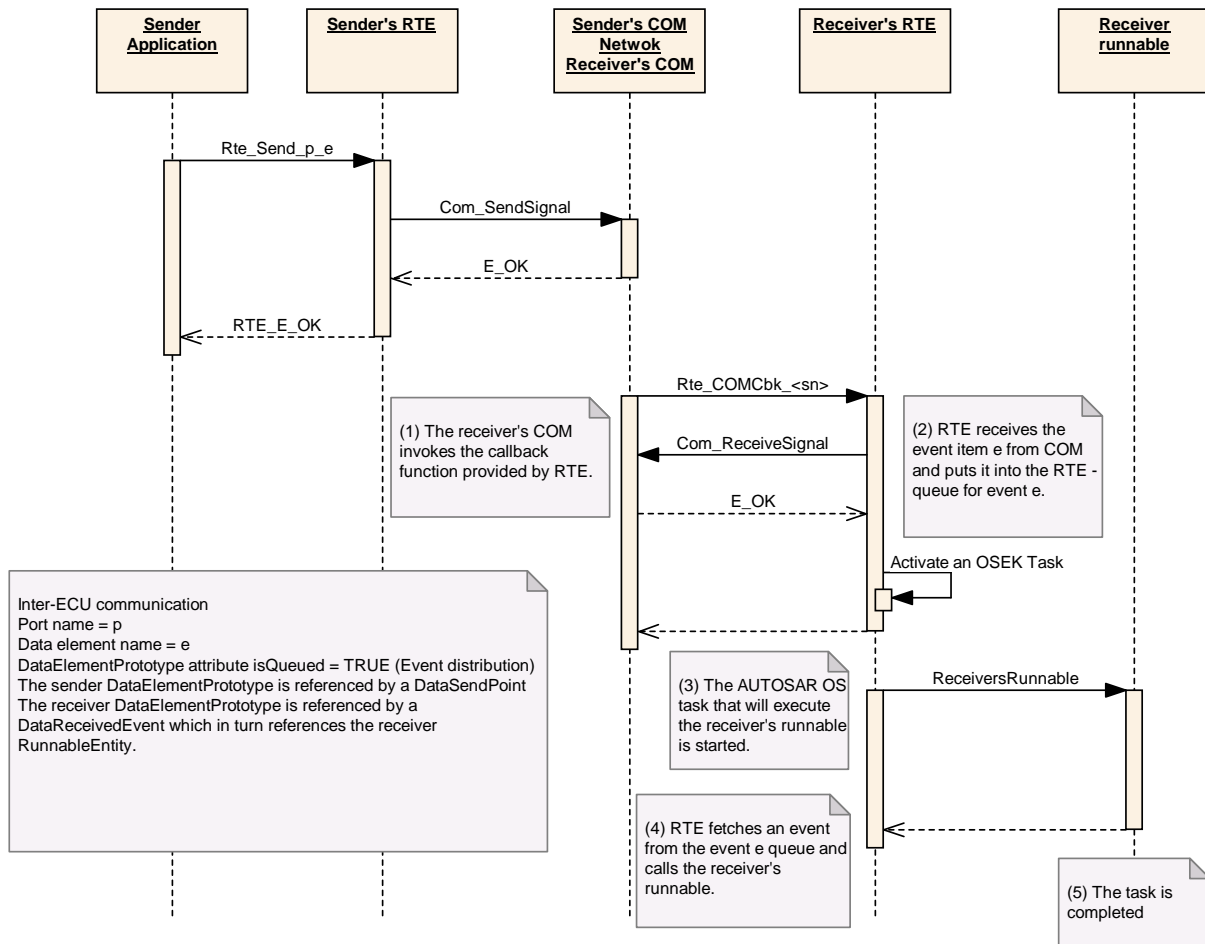


Figure 4.26: Sender Receiver communication with isQueued true and activation of runnable entity as reception mechanism

4.3.2 Client-Server

4.3.2.1 Introduction

Client-server communication involves two entities, the client which is the requirer (or user) of a service and the server that provides the service.

The client initiates the communication, requesting that the server performs a service, transferring a parameter set if necessary. The server, in the form of the RTE, waits for incoming communication requests from a client, performs the requested service and dispatches a response to the client's request. So, the direction of initiation is used to categorize whether a AUTOSAR software-component is a client or a server.

A single component can be both a client and a server depending on the software realization.

The invocation of a server is performed by the RTE itself when a request is made by a client. The invocation occurs synchronously with respect to the RTE (typically via a function call) however the client's invocation can be either synchronous (wait for server to complete) or asynchronous with respect to the server.

[rte_sws_6019] The only mechanism through which a server can be invoked is through a client-server invocation request from a client [RTE00029].

The above requirement means that *direct invocation* of the function implementing the server outside the scope of the RTE is not permitted.

A server has a dedicated provide port and a client has a dedicated require port. To be able to connect a client and a server, both ports must be categorized by the same interface.

The client can be blocked (synchronous communication) respectively non-blocked (asynchronous communication) after the service request is initiated until the response of the server is received.

A server implemented by a RunnableEntity with attribute `canBeInvokedConcurrently` set to `FALSE` is not allowed to be invoked concurrently and since a server can have one or more clients the server may have to handle concurrent service calls (n:1 communication) the RTE must ensure that concurrent calls do not interfere.

[rte_sws_4515] It is the responsibility of the RTE to ensure that serialization⁷ of the operation is enforced when the server runnable attribute `canBeInvokedConcurrently` is `FALSE`. Note that the same server may be called using both synchronous and asynchronous communication [RTE00033].

⁷Serialization ensures at most one thread of control is executing an instance of a runnable entity at any one time. An AUTOSAR software-component can have multiple instances (and therefore a runnable entity can also have multiple instances). Each instance represents a different server and can be executed in parallel by different threads of control thus serialization only applies to an individual instance of a runnable entity – multiple runnable entities within the same component instance may also be executed in parallel.

[rte_sws_4516] The RTE's implementation of the client-server communication has to ensure that a service result is dispatched to the correct client if more than one client uses a service [RTE00080].

The result of the client/server operation can be collected using “wake up of wait point”, “explicit data read access” or “activation of runnable entity”.

If the client and server are executing in the same ECU, i.e. intra ECU Client-Server communication, the RTE API call for client-server communication (see Sect. 5.6.10) can be optimized to a direct function call of the client without any interaction with the RTE or the communication service. Since the communication occurs conceptually via the RTE (it is initiated via an RTE API call) the optimization does not violate the requirement that servers are only invoked via client-server requests.

4.3.2.2 Multiplicity

Client-server interfaces contain two dimensions of multiplicity; multiple clients invoking a single server and multiple operations within a client-server interface.

4.3.2.2.1 Multiple Clients Single Server

Client-server communication involves an AUTOSAR software-component invoking a defined “server” operation in another AUTOSAR software-component which may or may not return a reply.

[rte_sws_4519] The RTE shall support multiple clients invoking the same server operation ('n:1' communication where $n \geq 1$). [RTE00029]

4.3.2.2.2 Multiple operations

A client-server interface contains one or more operations. A port of a AUTOSAR software-component that *requires* an AUTOSAR client-server interface to the component can independently invoke any of the operations defined in the interface [RTE00089].

[rte_sws_4517] The RTE API shall support independent access to operations in a client-server interface [RTE00029].

Example 4.3

Consider a client-server interface that has two operations, `op1` and `op2` and that an AUTOSAR software-component definition requires a port typed by the interface. As a result, the RTE generator will create two API calls; one to invoke `op1` and another to invoke `op2`. The calls can invoke the server

operations either synchronously or asynchronously depending on the configuration.

Recall that each data element in a sender-receiver interface is transmitted independently (see Section 4.3.1.3) and that the coherent transmission of multiple data items is achieved through combining multiple items into a single complex data type. The transmission of the parameters of an operation in a client-server interface is similar to a record since the RTE guarantees that all parameters are handled atomically [RTE00073].

[rte_sws_4518] The RTE shall treat the parameters (and results) of a client-server operation atomically [RTE00033].

However, unlike a sender-receiver interface, there is no facility to combine multiple client-server operations so that they are invoked as a group.

4.3.2.2.3 Single Client Multiple Server

The RTE is *not* required to support multiple server operations invoked by a single client component request ('1:n' communication where $n > 1$).

4.3.2.2.4 Serialization

Each client can invoke the server simultaneously and therefore the RTE is required to support multiple requests of servers. If the server requires serialization, the RTE has to ensure it.

[rte_sws_4520] The RTE shall support simultaneous invocation requests of a server operation. [RTE00080]

[rte_sws_4522] The RTE shall ensure that the runnable entity implementing a server operation has completed the processing of a request before it begins processing the next request, if serialization is required by the server operation, i.e. `canBeInvokedConcurrently` attribute set to `FALSE` [RTE00033].

When this requirement is met the operation is said to be “serialized”. A serialized server only accepts and processes requests atomically and thus avoids the potential for conflicting concurrent access.

Client requests that cannot be serviced immediately due to a server operation being “busy” are required to be queued pending processing. The presence and depth of the queue is configurable.

If the runnable entity implementing the server operation is reentrant, i.e. `canBeInvokedConcurrently` attribute set to `TRUE`, no serialization is necessary. This allows to implement invocations of reentrant server operations as direct function calls without involving the RTE.

4.3.2.3 Communication Time-out

The ServerCallPoint allows to specify a timeout so that the client can be notified that the server is not responding and can react accordingly. If the client invokes the server synchronously, the RTE API call to invoke the server reports the timeout. If the client invokes the server asynchronously, the timeout notification is passed to the client by the RTE as a return value of the API call that collects the result of the server operation.

[rte_sws_3763] The RTE shall ensure that timeout monitoring is performed for client-server communication, regardless of the receive mode for the result.

If the server is invoked asynchronously and a WaitPoint is specified to collect the result, two timeout values have to be specified, one for the ServerCallPoint and one for the WaitPoint.

[rte_sws_3764] If different timeout values are specified for the AsynchronousServerCallPoint and for the WaitPoint associated with the AsynchronousServerCallReturnsEvent for this AsynchronousServerCallPoint, the configuration shall be rejected by the RTE generator.

In asynchronous client-server communication the AsynchronousServerCallReturnsEvent associated with the AsynchronousServerCallPoint for an OperationPrototype shall indicate that the server communication is finished or that a timeout occurred. The status information about the success of the server operation shall be available as the return value of the RTE API call generated to collect the result.

[rte_sws_3765] For each asynchronous invocation of an operation prototype only one AsynchronousServerCallReturnsEvent shall be passed to the client component by the RTE. The AsynchronousServerCallReturnsEvent shall indicate either that the transmission was successful or that the transmission was not successful.

[rte_sws_3766] The status information about the success or failure of the asynchronous server invocation shall be available as the return value of the RTE API call to retrieve the result.

After a timeout was detected, no result shall be passed to the client.

[rte_sws_3770] If a timeout was detected by the RTE, no result shall be passed back to the client.

Since an asynchronous client can have only one outstanding server invocation at a time, the RTE has to monitor when the server can be safely invoked again. In normal operation, the server can be invoked again when the result of the previous invocation was collected by the client.

[rte_sws_3773] If a server is invoked asynchronously and no timeout occurred, the RTE shall ensure that the server can be invoked again by the same client, after the result was successfully passed to the client.

In intra-ECU client-server communication, the RTE can determine whether the server runnable is still running or not.

[rte_sws_3771] If a timeout was detected in asynchronous intra-ECU client-server communication, the RTE shall ensure that the server is not invoked again by the same client until the server runnable has terminated.

In inter-ECU communication, the client RTE has no knowledge about the actual status of the server. The response of the server could have been lost because of a communication error or because the server itself did not respond. Since the client-side RTE cannot distinguish the two cases, the client must be able to invoke the server again after a timeout expired.

[rte_sws_3772] If a timeout was detected in asynchronous inter-ECU client-server communication, the RTE shall ensure that the server can be invoked again by the same client after the timeout notification was passed to the client.

Note that this might lead to client and server running out of sync, i.e. the response of the server belongs to the previous, timed-out invocation of the client. The application has to handle the synchronization of client and server after a timeout occurred.

[rte_sws_3767] If the timeout value of the ServerCallPoint is 0, no timeout monitoring shall be performed.

If the `canBeInvokedConcurrently` attribute of the server runnable is set to `TRUE`, no timeout monitoring has to be performed to allow the optimization of the RTE API call to invoke the server to a direct function call.

[rte_sws_3768] If the `canBeInvokedConcurrently` attribute of the server runnable is set to `TRUE`, no timeout monitoring shall be performed if the RTE API call to invoke the server is implemented as a direct function call.

4.3.2.4 Port-Defined argument values

Port-defined argument values exist in order to support interaction between Application Software Components and Basic Software Modules.

Several Basic Software Modules use an integer identifier to represent an object that should be acted upon. For instance, the NVRAM Manager uses an integer identifier to represent the NVRAM block to access. This identifier is not known to the client, as the client must be location independent, and the NVRAM block to access for a given application software component cannot be identified until components have been mapped onto ECUs.

There is therefore a mismatch between the information available to the client and that required by the server. Port-defined argument values bridge that gap.

The required port-defined arguments (the fact that they are required, their data type and their values) are specified within the input to the RTE generator. (See requirements `rte_sws_in_1361` and `rte_sws_in_1362`.)

[rte_sws_1360] When invoking the runnable entity specified for an OperationInvokedEvent, the RTE must include the port-defined argument values between the instance handle (if it is included) and the operation-specific parameters, in the order they are given in the template.

Requirement `rte_sws_1360` means that a client will make a request for an operation on a require (Client-Server) port including only its instance handle (if required) and the explicit operation parameters, yet the server will be passed the implicit parameters as it requires.

Note that the values of implicit parameters are constant for a particular server runnable entity; it is therefore expected that using port-defined argument values imposes no RAM overhead (beyond any extra stack required to store the additional parameters).

4.3.2.5 Buffering

Client-Server-Communication is a two-way-communication. A request is sent from the client to the server and a response is sent back.

Unless a server call is implemented as direct function call, the RTE shall store or buffer the communication on the corresponding receiving sides, requests on server side and responses on client side, respectively:

- **[rte_sws_2527]** Unless a server call is implemented as a direct function call, the RTE shall buffer a request on the server side in a first-in-first-out queue as described in chapter 4.3.1.10.2 for queued data elements.
- **[rte_sws_2528]** Unless a server call is implemented as a direct function call, RTE shall keep the response on the client side in a queue with queue length 1.

For the server side, the attribute `queueLength` of the `ServerComSpec` specifies the length of the queue.

[rte_sws_2529] The RTE shall reject a queue of length ≤ 0 .

[rte_sws_2530] The RTE shall use the queue of requests to serialise access to a server.

A buffer overflow of the server is not reported to the client. The client will receive a time out.

[rte_sws_7008] If a server call is implemented by direct function call the RTE shall not create any copies for parameters passed by reference. Therefore, it is the responsibility of the application to provide consistency mechanisms for referenced parameters if necessary.

4.3.2.6 Inter ECU Response to Request Mapping

RTE is responsible to map a response to the corresponding request. With this mapping, RTE can activate or resume the corresponding runnable and provide the response to the correct client. The following situations can be distinguished:

- Mapping of a response to the correct request within one ECU. In general, this is solved already by the call stack. The details are implementation specific and will not be discussed in this document.
- Mapping of a response coming from a different ECU.

The problem of request to response mapping in inter-ECU communication can be split into:

- Mapping of a response to the correct client. This is discussed in 4.3.2.6.1.
- Mapping of a response to the correct request within of one client. This is discussed in 4.3.2.6.2.

The general approach for the inter-ECU request response mapping is to use transaction handles.

[rte_sws_2649] The transaction handle shall contain two parts of unsigned integer type with configurable size,

- the client identifier
- and a sequence counter.

[rte_sws_2651] The transaction handle shall be used for the identification of client server transactions communicated via COM.

[rte_sws_2652] The transaction handle shall be bundled with the parameters of a request or response in the same signal group.

[rte_sws_2653] The RTE on the server side shall return the transaction handle of the request without modification together with the response.

Since there is always at most one open request per client (see `rte_sws_2658`), the transaction handle can be kept within the RTE and does not have to be exposed to the SW-C.

4.3.2.6.1 Client Identity

The RTE uses the following mechanism to implement client identity:

[rte_sws_2579] In case of a server on one ECU with multiple clients on other ECUs, the client server communication shall use different unique COM signals and signal groups for each client to allow the identification of the client associated with each system signal.

With this mechanism, the server-side RTE must handle the fan-in. This is done in the same way as for sender-receiver communication.

[rte_sws_3769] If multiple clients have access to one server, the RTE on the server side has to queue all incoming server invocations while ensuring data consistency.

[rte_sws_5066] The data type used to hold the client identifier shall be derived from the system template's [14] **length** attribute of the corresponding **SystemSignal** referenced by the **ClientIdMapping**.

The structure is shown in figure 4.27.

4.3.2.6.2 SequenceCounter

The purpose of sequence counters is to map a response to the correct request of a known client.

[rte_sws_2658] In case of inter-ECU communication, RTE shall allow only one request per client and server operation at any time.

rte_sws_2658 does not apply to intra-ECU communication.

rte_sws_2658 implies under normal operation that a response can be mapped to the previous request. But, when a request or response is lost or delayed, this order can get out of phase. To allow a recovery from lost or delayed signals, a sequence counter is used. The sequence counter can also be used to detect stale responses after a restart of the client side RTE and SW-C.

[rte_sws_2654] RTE shall have a sequence counter for each inter ECU client server connection.

[rte_sws_2655] RTE shall initialize all sequence counters with zero during `Rte_Start`.

[rte_sws_2656] RTE shall increase each sequence counter in a cyclic manner after a client server operation has finished successfully or with a timeout.

[rte_sws_2657] RTE shall ignore incoming responses that do not match the sequence counter.

[rte_sws_5067] The data type used to hold the sequence counter shall be derived from the system template's [14] **length** attribute of the corresponding **SystemSignal** referenced by the **SequenceCounterMapping**.

The structure is shown in figure 4.27.

4.3.2.7 Operation

4.3.2.7.1 Inter-ECU Mapping

The client server protocol defines how a client call and the server response are mapped onto the communication infrastructure of AUTOSAR in case of inter-ECU communication. This allows RTE implementations from different vendors to interpret the client server communication in the same way.

The AUTOSAR System Template [14] does specify a protocol for the client server communication in AUTOSAR. A short overview of the major elements is provided in this section.

The structure in figure 4.27 describes the client server protocol as defined in the AUTOSAR System Template [14].

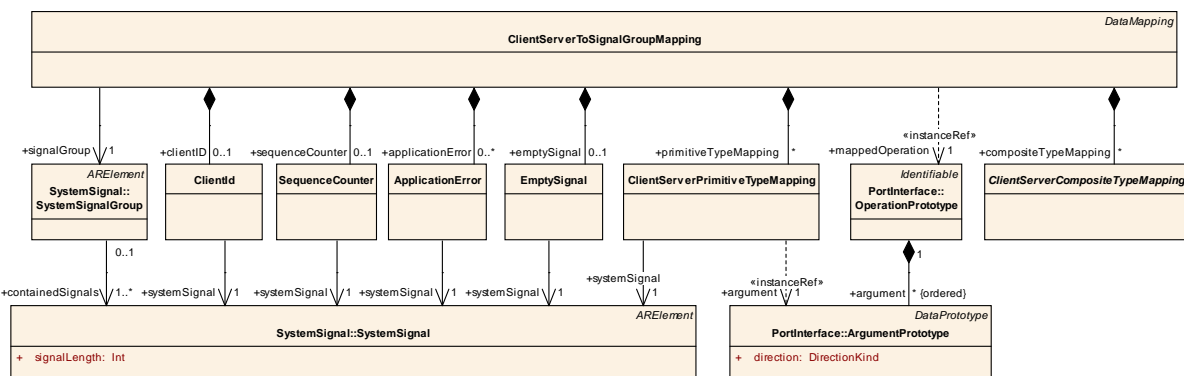


Figure 4.27: Standardized client server protocol

For each **OperationPrototype** defined at a **PortPrototype** two **ClientServerToSignalGroupMapping** objects have to be defined representing the server call and the response.

[rte_sws_5054] The RTE Generator shall reject an input configuration where for any **OperationPrototype** of any **PortPrototype** there are no two **ClientServerToSignalGroupMappings** defined, one representing the server call and the other representing the response.

[rte_sws_5055] The RTE Generator shall use the **ClientServerToSignalGroupMapping** information to establish the configuration with the lower layers of AUTOSAR (e.g. COM).

[rte_sws_6028] The arguments, application errors, client identifier, and sequence counter of an operation shall be mapped to two dedicated composite data items; one for the request and one for the response.

Each **ClientServerToSignalGroupMapping** references a unique **SystemSignalGroup** which holds all the signals related to the call or response.

For each **ArgumentPrototype** either a **ClientServerPrimitiveTypeMapping** or a **ClientServerCompositeTypeMapping** is defined which maps the operation arguments to **SystemSignal** elements.

[rte_sws_5056] If a **ClientId** element is configured it references the **SystemSignal** which holds the Client Id (see section 4.3.2.6.1). The RTE Generator shall utilize this **SystemSignal** as the **ClientId**.

[rte_sws_5057] If a **SequenceCounter** element is configured it references the **SystemSignal** which holds the Sequence Counter (see section 4.3.2.6.2). The RTE Generator shall utilize this **SystemSignal** as the **SequenceCounter**.

[rte_sws_5058] If an **ApplicationError** element is configured it references the **SystemSignal** which holds the Application Error (see section 5.2.6.6). The RTE Generator shall utilize this **SystemSignal** as the **ApplicationError**.

There might be configuration where no actual data is transferred between the client and the server (or vice versa). In this case a **SystemSignalGroup** shall be used with an update bit defined in System Description. In this case at least one **SystemSignal** is required to be present in the **SystemSignalGroup**.

[rte_sws_5059] If no actual data is configured for a client server communication the element **EmptySignal** shall reference a zero length **SystemSignal**. In this case the RTE shall send the **SignalGroup** to initialize the communication.

4.3.2.7.2 Atomicity

The requirements for atomicity from Section 4.3.1.11.2 also apply for the composite data types described in Section 4.3.2.7.1.

4.3.2.7.3 Fault detection and reporting

Client Server communication may encounter interruption like:

- Buffer overflow at the server side.
- Communication interruption.
- Server might be inaccessible for some reason.

The client specifies a timeout that will expire in case the server or communication fails to complete within the specified time. The reporting method of an expired timeout depends on the communication attributes:

- If the C/S communication is synchronous the RTE returns `RTE_E_TIMEOUT` on the `Rte_Call` function (see chapter 5.6.10).
- If the C/S communication is asynchronous the RTE returns `RTE_E_TIMEOUT` on the `Rte_Result` function (see chapter 5.6.11).

In the case that RTE detects that the COM service is not available when forwarding signals to COM, the RTE returns `RTE_E_COM_STOPPED` on the `Rte_Call` (see chapter 5.6.10).

If the client still has an outstanding server invocation when the server is invoked again, the RTE returns `RTE_E_LIMIT` on the `Rte_Call` (see chapter 5.6.10).

In the absence of structural errors, application errors will be reported if present.

4.3.2.7.4 Asynchronous Client Server communication

Figure 4.28 shows a sequence diagram of how asynchronous client server communication may be implemented by RTE.

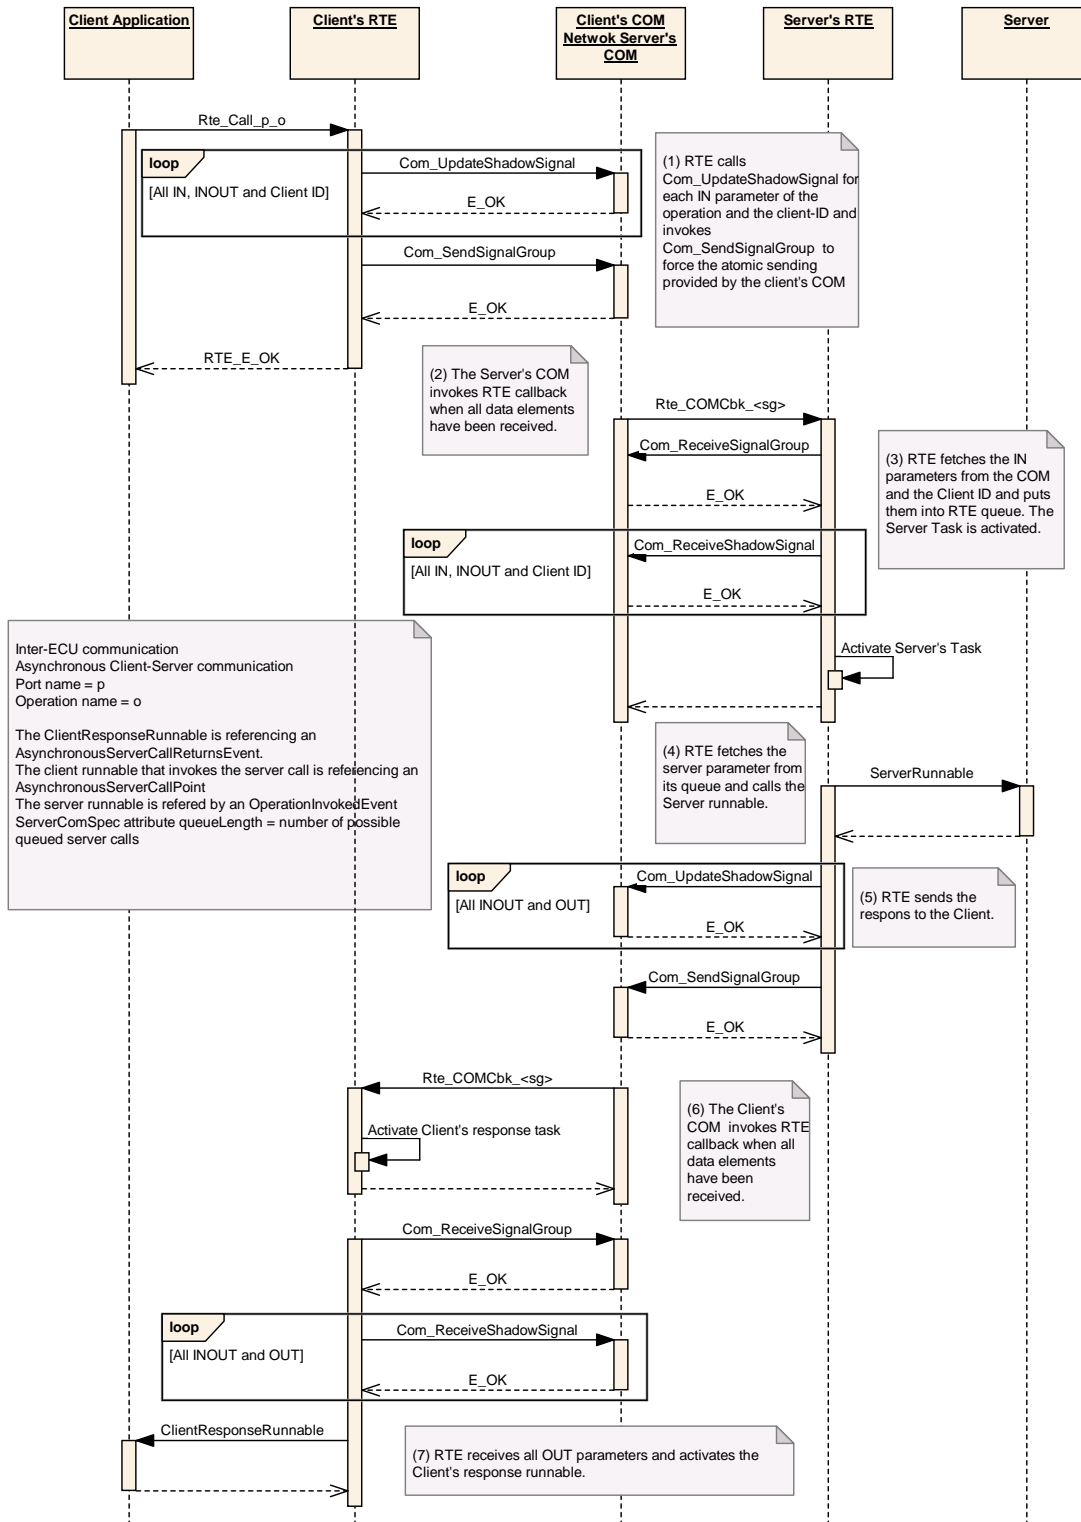


Figure 4.28: Client Server asynchronous

4.3.2.7.5 Synchronous Client Server communication

Figure 4.29 shows a sequence diagram of how synchronous client server communication may be implemented by RTE.

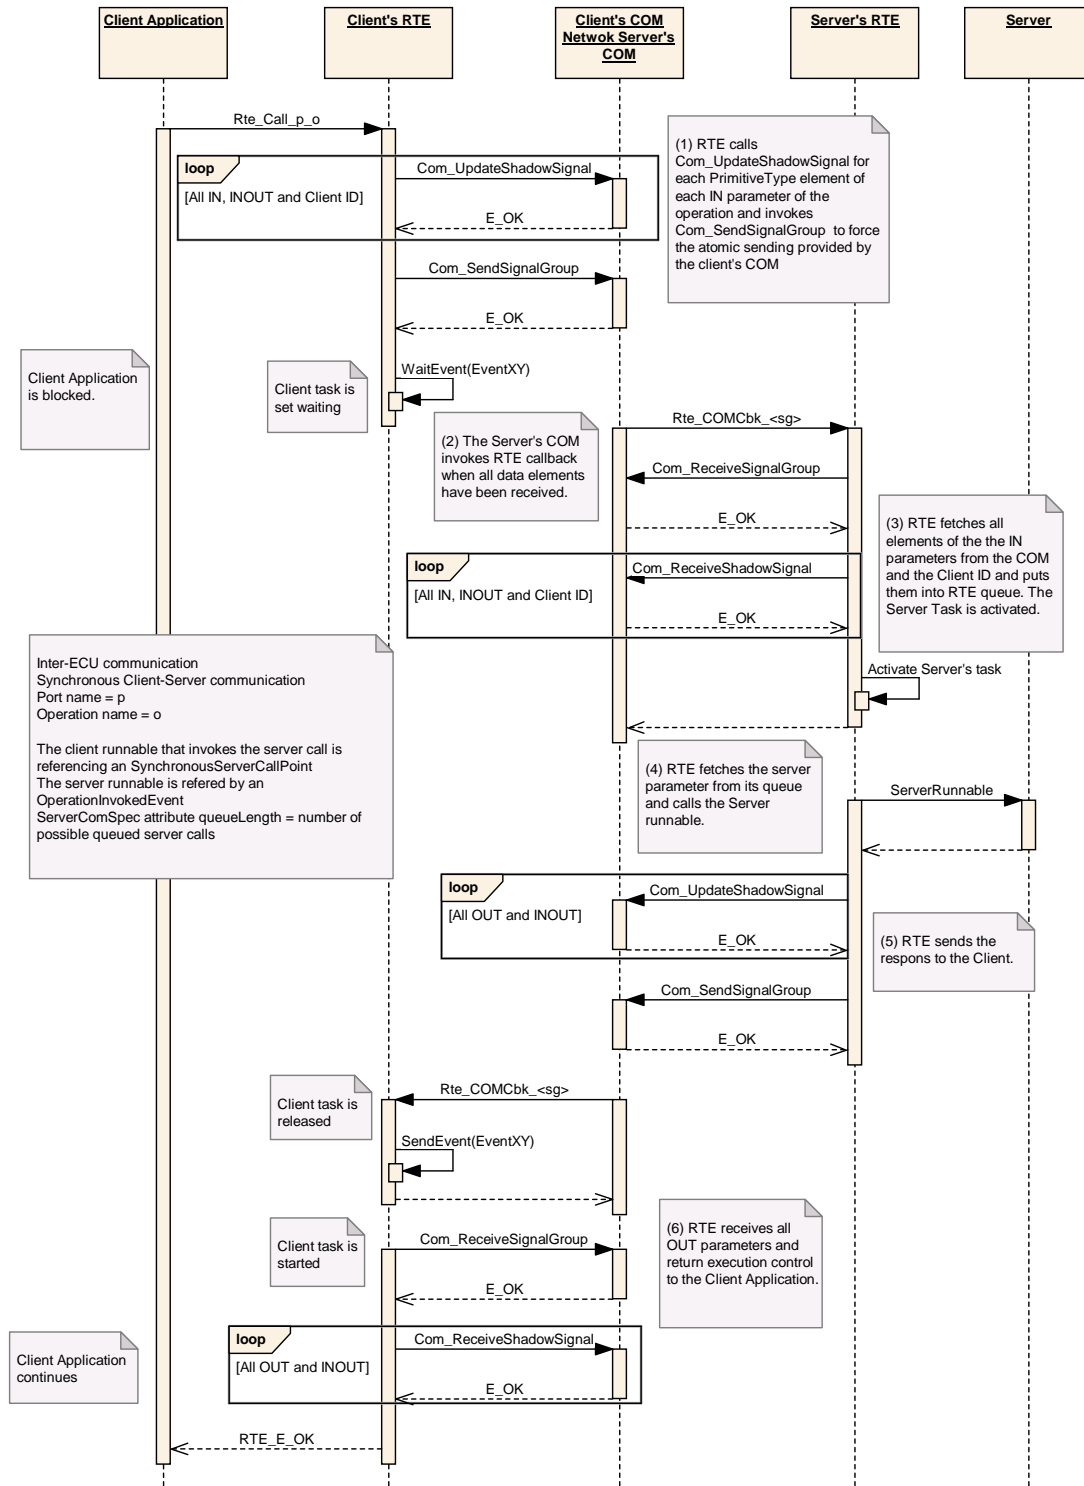


Figure 4.29: Client Server synchronous

4.3.3 SWC internal communication

4.3.3.1 InterRunnableVariables

Sender/Receiver and Client/Server communication through AUTOSAR ports are the model for communication between AUTOSAR SW-Cs.

For communication between Runnables inside of an AUTOSAR SW-C the AUTOSAR SW-C Template [17] establishes a separate mechanism. Non-composite AUTOSAR SW-C can reserve InterRunnableVariables which can only be accessed by the Runnables of this one AUTOSAR SW-C. The Runnables might be running in the same or in different task contexts. Read and write accesses are possible.

[rte_sws_3589] The RTE has to support InterRunnableVariables for single and multiple instances of AUTOSAR SW-Cs.

InterRunnableVariables have a behavior corresponding to Sender/Receiver communication *between* AUTOSAR SW-Cs (or rather between Runnables of different AUTOSAR SW-Cs).

But why not use Sender/Receiver communication directly instead? Purpose is data encapsulation / data hiding. Access to InterRunnableVariables of an AUTOSAR SW-C from other AUTOSAR SW-Cs is not possible and not supported by RTE. InterRunnableVariable content stays SW-C internal and so no other SW-C can use. Especially not misuse it without understanding how the data behaves.

Like in Sender/Receiver (S/R) communication between AUTOSAR SW-Cs two different behaviors exist:

1. InterRunnableVariables with **implicit** behavior

This behavior corresponds with *DataReadAccess* / *DataWriteAccess* of Sender/Receiver communication and is supported by *implicit S/R API* in this specification.

Note:

If *DataWriteAccess* is specified for a certain interrunnable variable, but no RTE API for implicit write of this interrunnable variable is called during an execution of the runnable, an undefined value is written back when the runnable terminates.

For more details see section 4.2.4.6.1.

For APIs see sections 5.6.20 and 5.6.21.

2. InterRunnableVariables with **explicit** behavior

This behavior corresponds with *DataSendPoint* / *DataReceivePoint* of Sender/Receiver communication and is supported by *explicit S/R API* in this specification.

For more details see section 4.2.4.6.2

For APIs see sections 5.6.22 and 5.6.23.

4.4 Modes

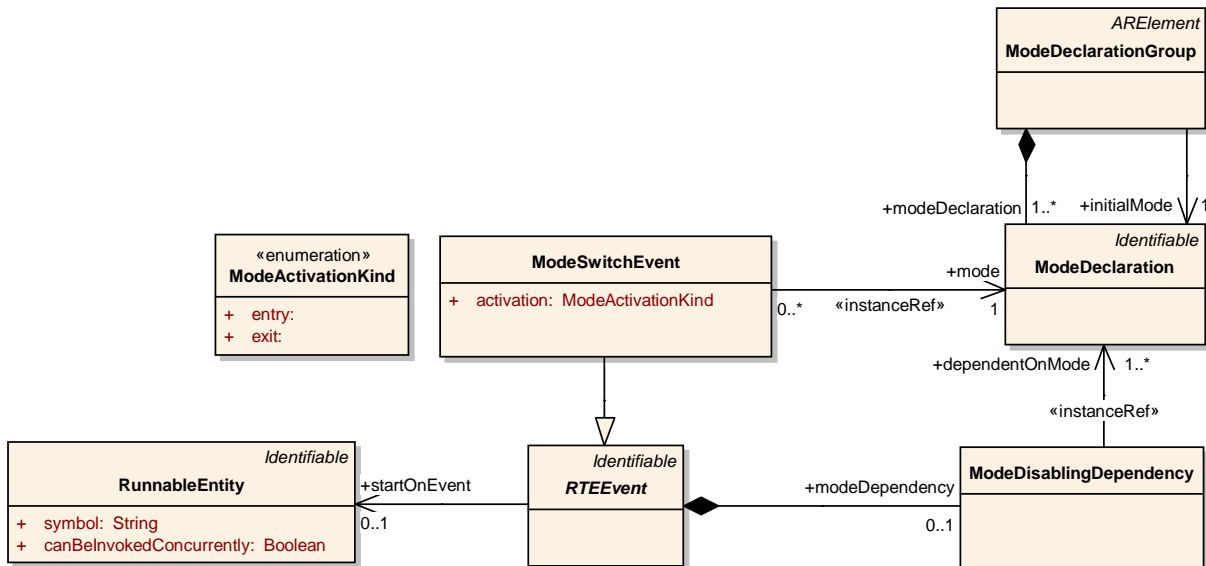


Figure 4.30: Summary of the use of ModeDeclarations by an AUTOSAR Software-Component instance as defined in the Software Component Template Specification [17].

The purpose of modes is to start runnables on the transition between modes and to disable (/enable) specified triggers of runnables in certain modes. Here, we use the specification of modes from the Software Component Template Specification [17].

The first subsection 4.4.1 describes how modes can be used by an AUTOSAR software-component `mode user()`. The role of the `mode manager` who initiates mode switches is described in section 4.4.2. How ModeDeclarations are connected to a state machine is described in subsection 4.4.3. The behaviour of the RTE regarding mode switches is detailed in subsection 4.4.4.

One usecase of modes is described in section 4.5.2 for the initialization and finalization of AUTOSAR Software-Components. Modes can be used for handling of communication states as well as for specific application purposes. The specific definition of modes and their use is not in the scope of this document.

The status of the modes will be notified to the `mode user` by a specific form of sender receiver communication - `mode switch notifications` - as described in the subsection 4.4.5. The port for receiving (or sending) a `mode switch notification` is called `mode port`.

4.4.1 Mode User

To use modes, an AUTOSAR software-component (`mode user`) has to reference a `ModeDeclarationGroup` by a `ModeDeclarationGroupPrototype` of a `require mode port`, see section 4.4.5. The `ModeDeclarationGroup` contains the required modes.

The ModeDeclarations can be used in two ways by the `mode user` (see also figure 4.30):

1. Modes can be used to trigger runnables: The InternalBehavior of the AUTOSAR SW-C can define a ModeSwitchEvent referencing the required ModeDeclaration. This ModeSwitchEvent can then be used as trigger for a runnable. The ModeSwitchEvent carries an attribute ModeActivationKind which can be 'exit' or 'entry'.

A runnable that is triggered by a ModeSwitchEvent with ModeActivationKind 'exit' is triggered on exiting the mode. For simplicity it will be called `OnExit runnable`. Correspondingly, an `OnEntry runnable` is triggered by a ModeSwitchEvent with ModeActivationKind 'entry' and will be executed when the mode is entered.

Since a runnable can be triggered by multiple RteEvents, it can be both, an OnExit- and OnEntry runnable

RTE does not support a wait point for a ModeSwitchEvent (see `rte_sws_1358`).

2. An RTEEvent that starts a runnable can contain a ModeDisablingDependency which references a ModeDeclaration.

[rte_sws_2503] If a runnable entity *r* is referenced with `startOnEvent` by an RTE-Event *e* that has a ModeDisablingDependency on a mode *m*, then RTE shall not `activate runnable r` on any occurrence of *e* while the mode *m* is active.

Note: As a consequence of `rte_sws_2503` in combination with `rte_sws_2661`, RTE will not `start runnable r` on any occurrence of *e* while the mode *m* is active.

The `mode disabling` is active during the transition to a mode, during the mode itself and during the transition for exiting the mode. For a precise definition see section 4.4.4.

The existence of a ModeDisablingDependency prevents the RTE to start the `mode disabling dependent runnable` by the disabled RTEEvent during the mode, referenced by the ModeDisablingDependency, and during the transitions from and to that mode. ModeDisablingDependencies override any activation of a runnable by the disabled RTEEvents. This is also true for the ModeSwitchEvent.

A runnable can not be 'enabled' explicitly. A runnable is only 'enabled' by the absence of any active ModeDisablingDependency.

Note that ModeDisablingDependencies do not prevent the wake up from a Wait-Point by the 'disabled' RTEEvent.

[rte_sws_2504] The existence of a ModeDisablingDependency shall not instruct the RTE to kill or preempt a running runnable at a mode switch.

The RTE might switch schedule tables to implement mode disabling dependencies for cyclic triggers of runnables.

- To do this, the RTE generator needs to know mutual exclusivity and coverage of modes, see `rte_sws_2542`.
- **[`rte_sws_ext_2559`]** The RTE configurator shall have access to the schedule table configuration (see also `rte_sws_4014`).

4.4.2 Mode Manager

Entering and leaving modes is initiated by a `mode manager`. A `mode manager` might be a basic software module, for example the communication manager (COMM) or the ECU state manager. The `mode manager` may also be an AUTOSAR SW-C. In this case, it is called an `application mode manager`. The modes of an `application mode manager` do not have to be standardized.

The `mode manager` contains the master state machine to represent the modes. The `mode manager` has a `provide mode port` to communicate the current mode to the `mode users` via the RTE.

The RTE will take the actions necessary to switch between the modes. This includes the termination and execution of several `runnables` from all `mode users` that are connected to the same `ModeDeclarationGroupPrototype` of the `mode manager`. To do so, the RTE needs a state machine to keep track of the currently active modes and transitions initiated by the `mode manager`. The RTE's mode machine is called `mode machine instance`. There is exactly one `mode machine instance` for each `ModeDeclarationGroupPrototype` of a `mode manager's provide mode port`.

It is the responsibility of the `mode manager` to advance the RTE's `mode machine instance` by sending `mode switch notifications` to the `mode users`. The `mode switch notifications` are implemented by a non blocking API (see 5.6.5). So, the `mode switch notifications` alone provide only a loose coupling between the state machine of the `mode manager` and the `mode machine instance` of the RTE. To prevent, that the `mode machine instance` lags behind and the states of the `mode manager` and the RTE get out of phase, the `mode manager` can use `acknowledgment feedback` for the `mode switch notification`. RTE can be configured to send an `acknowledgment` of the `mode switch notification` to the `mode manager` when the requested transition is completed.

At the `mode manager`, the `acknowledgment` results in an `ModeSwitchedAckEvent`. As with `DataSendCompletedEvents`, this event can be picked up with the polling or blocking `Rte_Feedback` API. And the event can be used to trigger `runnables` to pick up the status.

Some possible usage patterns for the `acknowledgment` are:

- The most straight forward method is to use a sequence of `Rte_Switch` and a blocking `Rte_Feedback` to send the `mode switch notification` and wait for the completion. This requires the use of an extended task.

- Another possibility is to have a cyclic runnable (maybe the same that switches the modes) to poll for the feedback.
- The feedback can also be polled from a runnable that is started by the `ModeSwitchedAckEvent`.

The `mode manager` can also use the `Rte_Mode` API to read the currently active mode from the RTE's perspective.

4.4.3 Refinement of the semantics of ModeDeclarations and ModeDeclaration-Groups

To implement the logic of mode switches, the RTE needs some basic information about the available modes. For this reason, RTE will make the following additional assumptions about the modes of one `ModeDeclarationGroup`:

1. **[rte_sws_ext_2542]** Whenever any runnable entity is running, there shall always be exactly one mode or one mode transition active of each `ModeDeclarationGroup`.
2. Immediately after initialization, RTE will execute a transition to the initial mode of each `ModeDeclarationGroup` (see `rte_sws_2544`).

RTE will enforce the `mode disabling`s of the initial modes and trigger the `OnEntry` `runnables` of the initial modes (if there are any) immediately after initialization.

In other words, RTE assumes, that the modes of one `ModeDeclarationGroup` belong to exactly one state machine without nested states. The state machines cover the whole lifetime⁸ of the atomic AUTOSAR SW-Cs.

4.4.4 Order of actions taken by the RTE upon interception of a mode switch notification

This section describes what the 'communication' of a mode switch to a `mode user` actually does. What does the RTE do to switch a mode and especially in which order.

Typical usage of modes to protect resources

RTE can start the execution of runnables and prevent the execution of runnables. In the context of mode switches,

- RTE starts `OnExit` `runnables` for leaving the previous mode. This is typically used by 'clean up runnables' to free resources that were used during the previous mode.

⁸The lifetime of an atomic AUTOSAR SW-C is considered to be the time span in which the SW-C's runnables are being executed.

- RTE starts `OnEntry` runnables for entering the next mode. This is typically used by 'initialization runnables' to allocate resources that are used in the next mode.
- And RTE can prevent the execution of mode disabling dependent runnables within a mode. This is typically used with time triggered 'work runnables' that use a resource which is not available in a certain mode.

According to this use case, during the execution of 'clean up runnables' and 'initialization runnables' the 'work runnables' should be disabled to protect the resource. Also, if the same resource is used (by different SW-C's) in two successive modes, the 'clean up runnables' should be safely terminated before the 'initialization runnables' of the next mode are executed. In summary, this would lead to the following sequence of actions by the RTE upon reception of the `mode switch` notification:

1. activate `mode disables` for the next mode
2. wait for the newly disabled runnables to terminate
3. execute 'clean up runnables'
4. wait for the 'clean up runnables' to terminate
5. execute 'initialization runnables'
6. wait for the 'initialization runnables' to terminate
7. deactivate `mode disables` for the previous modes and enable runnables that have been disabled in the previous mode.

Often, only a fraction of the SW-Cs and runnables of one ECU depends on the modes that are switched. Consequently, it should be possible to design the system in a way, that the mode switch does not influence the performance of the remaining software.

The remainder of this section lists the requirements that guarantee the behavior described above.

All runnables with dependencies on modes have to be executed or terminated during mode transitions. Restriction `rte_sws_2500` requires these runnables to be of category 1 to guarantee finite execution time.

For simplicity of the implementation to guarantee the order of runnable executions, the following restriction is made:

All `OnEntry` runnables and `OnExit` runnables of the same mode machine instance should be mapped to the same task (see `rte_sws_2662`).

[`rte_sws_2667`] Within the mode manager's `Rte_Switch` API call to indicate a mode switch, one of the following shall be done:

1. If the corresponding mode machine instance is in a transition, and the queue for mode switch notifications is full, `Rte_Switch` shall return an error immediately.

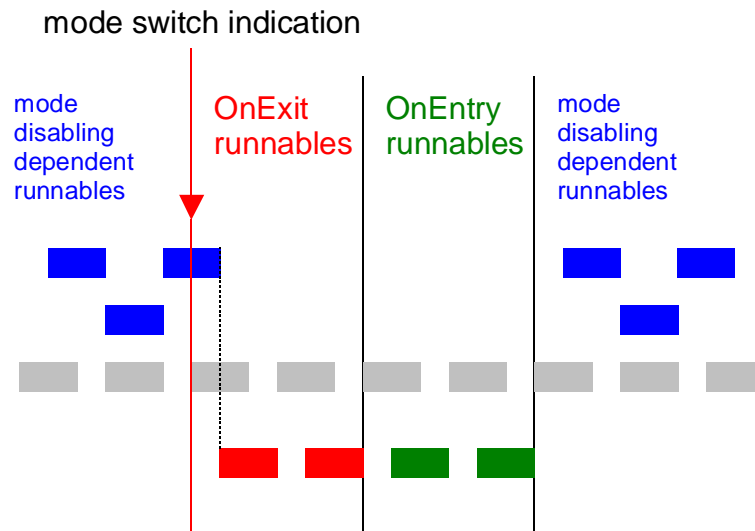


Figure 4.31: This figure shall illustrate what kind of runnables will run in what order during a mode transition. The boxes indicate activated runnables. Mode disabling dependant runnables are printed in blue. OnExit and OnEntry runnables are printed in red and green.

2. If the corresponding mode machine instance is in a transition, and the queue for mode switch notifications is not full, the mode switch notification shall be queued.
3. If the mode machine instance is not in a transition, `RteSwitch` shall activate the mode disabling (see `rte_sws_2661`) of the next mode, and initiate the transition as described by the sequence in `rte_sws_2665`.

The following list holds the requirements for the steps of a mode transition.

- **[rte_sws_2661]** At the beginning of a transition of a mode machine instance, the RTE shall activate the mode disabling of the next mode (see also `rte_sws_2503`), if any `ModeDisablingDependencies` for that mode are defined and wait for the termination of the newly disabled runnables.

Note: To guarantee that all activated mode disabling dependent runnables of this mode machine instance have terminated before the start of the `OnExit` runnables of the transition, RTE can exploit the restriction `rte_sws_2663` that mode disabling dependent runnables run with higher or equal priority than the `OnExit` runnables and the `OnEntry` runnables.

- **[rte_sws_2562]** RTE shall execute (activate and wait for termination) the `OnExit` runnables of the previous mode.
- **[rte_sws_2564]** RTE shall execute (activate and wait for termination) the `OnEntry` runnables of the next mode.

The `OnExit` runnables of the previous mode are executed before the `OnEntry` runnables of the next mode within the same task, see `rte_sws_2664` and `rte_sws_2662`.

- **[rte_sws_2563]** The RTE shall deactivate the previous `mode_disablings` and only keep the `mode_disablings` of the next mode.

With this, the transition is completed.

- **[rte_sws_2587]** At the end of the transition, RTE shall trigger the `ModeSwitchedAckEvents` connected to the `mode manager's ModeDeclarationGroupPrototype`.

This will result in an acknowledgment on the `mode manager's` side which allows the `mode manager` to wait for the completion of the mode switch.

[rte_sws_2665] During a transition of a `mode machine instance` each applicable of the steps

1. `rte_sws_2661` (The transition is entered in parallel with this step),
2. `rte_sws_2562`,
3. `rte_sws_2564`,
4. `rte_sws_2563` (The transition is completed with this step), and
5. immediately followed by `rte_sws_2587`

shall be executed in the order as listed. If a step is not applicable, the order of the remaining steps shall be unchanged.

[rte_sws_2668] Immediately after the execution of a transition as described in `rte_sws_2665`, RTE shall check the queue for pending `mode switch notifications` of this `mode machine instance`. If a `mode switch notification` can be dequeued, the `mode machine instance` shall enter the corresponding transition directly as described by the sequence in `rte_sws_2665`.

The `Rte_Mode` API will not indicate an intermediate mode, if a `mode switch notification` to the next mode is indicated before the transition to the intermediate mode is completed.

[rte_sws_2544] RTE shall initiate the transition to the initial modes of each `mode machine instance` during `Rte_Start`. During the transition to the initial modes, the steps defined in the following requirements have to be omitted as no previous mode is defined:

- `rte_sws_2562`
- `rte_sws_2563`
- `rte_sws_2587`

If applicable, the steps described by the following requirements still have to be executed for entering the initial mode:

- `rte_sws_2661`
- `rte_sws_2564`

[rte_sws_2630] RTE shall execute all steps of a mode switch (see `rte_sws_2661`, `rte_sws_2562`, `rte_sws_2563`, `rte_sws_2564`, and `rte_sws_2587`) synchronously for the whole `mode machine instance`.

I.e., the mode transitions will be executed synchronously for all `mode users` that are connected to the same `mode manager`'s `ModeDeclarationGroupPrototype`.

[rte_sws_2669] If the next mode and the previous mode of a transition are the same, the transition shall still be executed.

4.4.5 Notification of mode switches

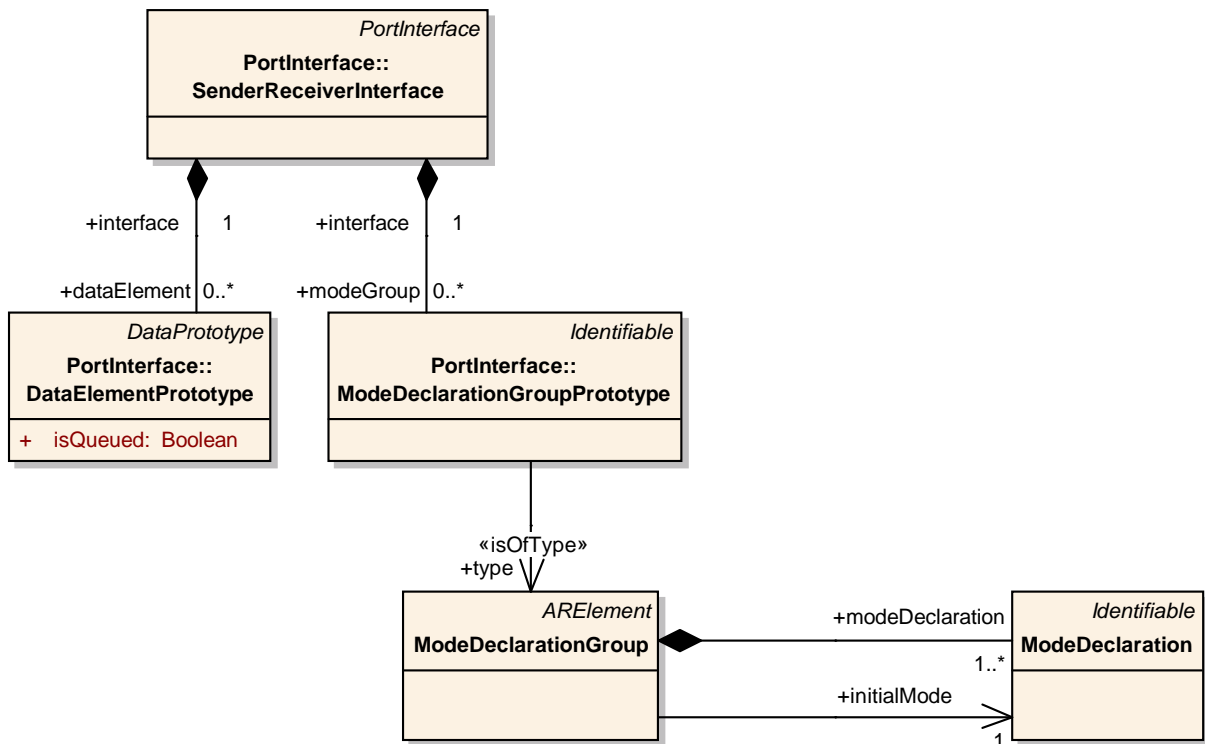


Figure 4.32: Definition of a mode switch interface.

- **[rte_sws_2549]** Mode switches shall be communicated by `ModeDeclarationGroupPrototypes` of a `SenderReceiverInterface` as defined in [17], see Fig. 4.32.

A `SenderReceiverInterface` with a `ModeDeclarationGroupPrototype` of a `ModeDeclarationGroup` will be called `mode switch interface` for the `ModeDeclarationGroup` in the following. The `mode ports` of the `mode manager` and the `mode user` are of the type of a `mode switch interface`.

- RTE only requires the notification of switches between modes.
- **[rte_sws_ext_2507]** The mode switch shall be notified to the mode user (and RTE) locally on each ECU.

[rte_sws_2586] The RTE generator shall reject a configuration with a nonlocal connection of a ModeDeclarationGroupPrototype.

Rationale: Even without communication to other ECUs, each state machine has to be in a well defined state/mode. This Requirement `rte_sws_ext_2507` does not prevent distributed mode management. But, for distributed mode management, a local agent is required on each ECU.

This implies that the connector between an `application mode manager` instance and the `mode user` instance can only be created after mapping of the SW-C to an ECU, because the `application mode manager` instance needs to be a specific agent, bound to one ECU.

- **[rte_sws_2508]** A mode switch shall be notified asynchronously as indicated by the use of a SenderReceiverInterface.

Rationale: This simplifies the communication. Due to `rte_sws_ext_2507` the communication is local and no handshake is required to guarantee reliable transmission.

RTE offers the api `Rte.Switch` to the mode manager for this notification, see 5.6.5.

- The `mode manager` might still require a feedback to keep it's internal state machine synchronized with the RTE view of active modes.

The RTE generator shall support an AcknowledgementRequest from the `mode port` of a `mode manager`, see `rte_sws_2587`, to notify the `mode manager` of the completion of a mode switch.

- **[rte_sws_2566]** A `mode switch interface` shall support 1:n communication.

Rationale: This simplifies the configuration and the communication. One mode switch can be notified to all receivers simultaneously.

A `mode switch interface` does not support n:1 communication, see `rte_sws_2670`.

- **[rte_sws_2624]** A mode switch shall be notified with event semantics, i.e., the mode switch notifications shall be buffered by RTE.

The queueing of mode switches (and ModeSwitchEvents) depends like that of DataReceivedEvents on the settings for the receiving port, see section 4.3.1.10.2.

- **[rte_sws_2567]** A `mode switch interface` shall only indicate the next mode of the transition.

The API takes a single parameter (plus, optionally, the instance handle) that indicates the requested 'next mode'. For this purpose, RTE will use identifiers of the modes as defined in `rte_sws_2568`.

- **[rte_sws_2546]** The RTE shall keep track of the active modes of a `mode manager's ModeDeclarationGroupPrototypes` (`mode machine instances`).

Rationale: This allows the RTE to guarantee consistency between the timing for firing of ModeSwitchEvents and disabling the start of runnables by ModeDisablingDependency without adding additional interfaces to a mode manager with fine grained substates on the transitions.

- The RTE offers an `Rte_Mode` API to the SW-C to get information about the active mode, see section 5.6.26.
- In addition to the `mode ports`, the `mode manager` may offer an AUTOSAR interface for requesting and releasing modes as a means to keep modes alive like for COMM and ECU State Manager.

4.5 Initialization and Finalization

4.5.1 Initialization and Finalization of the RTE

The ECU state manager calls the startup routine `Rte_Start` of the RTE at the end of startup phase II when the OS is available and all basic software modules are initialized.

[rte_sws_2513] The initialization routine of the RTE shall return within finite execution time.

Before the RTE is initialized completely, there is only a limited capability of RTE to handle incoming data from COM:

[rte_sws_2535] RTE shall ignore incoming client server communication requests, before RTE is initialized completely.

[rte_sws_2536] Incoming data and events from sender receiver communication shall be ignored, before RTE is initialized completely.

RTE will activate the `mode_disablings` of all initial modes during `RTEStart` and trigger the execution of the `OnEntry` `runnables` of the initial modes, see `rte_sws_2544`.

The finalization routine `RTE_Stop` of the RTE is called by the ECU state manager at the beginning of shutdown phase I when the OS is still available. (For details of the ECU state manager, see [8]. For details of `Rte_Start` and `Rte_Stop` see section 5.8.)

[rte_sws_2538] `Rte_Stop` shall stop the execution of all runnables.

4.5.2 Initialization and Finalization of AUTOSAR Software-Components

For the initialization and finalization of AUTOSAR software components, RTE provides the mechanism of mode switches. A `ModeSwitchEvent` of an appropriate `ModeDeclaration` can be used to trigger a corresponding initialization or finalization runnable (see `rte_sws_2562`). Runnables that shall not run during initialization or finalization can be disabled in the corresponding modes with a `ModeDisablingDependency` (see `rte_sws_2503`).

Since category 2 runnables have no predictable execution time and can not be terminated using `ModeDisablingDependencies`, it is the responsibility of the implementer to set meaningful termination criteria for the cat 2 runnables. These criteria could include mode information. At latest, all runnables will be terminated by RTE during the shutdown of RTE, see `rte_sws_2538`.

It is appropriate to use user defined modes that will be handled in a proprietary `application mode manager`.

All runnables that are triggered by entering an initial mode, are activated immediately after the initialization of RTE. They can be used for initialization. In many cases it might

be preferable to have a multi step initialization supported by a sequence of different initialization modes.

4.6 RTE Functionality Levels

There is a single RTE functionality level. So RTE is compliant AUTOSAR Functionality Conformance Class 1 (FCC1)

5 RTE Reference

“Everything should be as simple as possible, but no simpler.”

– *Albert Einstein*

5.1 Scope

This chapter presents the RTE API from the perspective of AUTOSAR applications and basic software – the same API applies to all software whether they are AUTOSAR software-components or basic software.

Section 5.2 presents basic principles of the API including naming conventions and supported programming languages. Section 5.3 describes the header files used by the RTE and the files created by an RTE generator. The data types used by the API are described in Section 5.5 and Sections 5.6 and 5.7 provide a reference to the RTE API itself including the definition of runnable entities. Section 5.10 defines the events that can be monitored during VFB tracing.

5.1.1 Programming Languages

The RTE is required to support components written using the C and C++ programming languages [RTE00126] as well as legacy software modules [RTE.IN016]. The ability for multiple languages to use the same generated RTE is an important step in reducing the complexity of RTE generation and therefore the scope for errors.

[rte_sws_1167] The RTE shall be generated in C.

[rte_sws_1168] All RTE code, whether generated or not, shall conform to the HIS subset of the MISRA C standard [18]. In technically reasonable, exceptional cases MISRA violations are permissible. Such violations shall be clearly identified and documented.

Specified MISRA violations are defined in Appendix D.

The RTE API presented in Section 5.6 is described using C. The API is also directly accessible from an AUTOSAR software-component written using C++ provided all API functions and instances of data structures are imported with C linkage.

[rte_sws_1011] The RTE generator shall ensure that, for a component written in C++, all imported RTE symbols are declared using C linkage.

For the RTE API for C and C++ components the import of symbols occurs within the application header file (Section 5.3.3).

5.1.2 Generator Principles

5.1.2.1 Operating Modes

An object-code component is compiled against an application header file that is created during the first “RTE Contract” phase of RTE generation. The object code is then linked against an RTE created during the second “RTE Generation” phase. To ensure that the object-code component and the RTE code are compatible the RTE generator supports *compatibility mode* that uses well-defined data structures and types for the component data structure. In addition, an RTE generator may support a *vendor* operating mode that removes compatibility between RTE generators from different vendors but permits implementation specific, and hence potentially more efficient, data structures and types.

[rte_sws_1195] All RTE operating modes shall be source-code compatible at the SW-C level.

Requirement `rte_sws_1195` ensures that a SW-C can be used in any operating mode as long as the source is available. The converse is not true – for example, an object-code SW-C compiled after the “RTE Contract” phase must be linked against an RTE created by an RTE generator operating in the same operating mode. If the vendor mode is used in the “RTE Contract” phase, an RTE generator from the same vendor (or one compatible to the vendor-mode features of the RTE generator used in the “RTE Contract” phase) has to be used for the “RTE Generation” phase.

5.1.2.1.1 Compatibility Mode

Compatibility mode is the default operating mode for an RTE generator and guarantees compatibility even between RTE generators from different vendors through the use of well-defined, “standardized”, data structures. The data structures that are used by the generated RTE in the compatibility mode are defined in Section 5.4.

Support for compatibility mode is required and therefore is guaranteed to be implemented by all RTE generators.

[rte_sws_1151] The *compatibility mode* shall be the default operating mode and shall be supported by all RTE generators, whether they are for the “RTE Contract” or “RTE Generation” phases.

The compatibility mode uses custom (generated) functions with standardized names and data structures that are defined during the “RTE Contract” phase and used when compiling object-code components.

[rte_sws_1216] SW-Cs that are compiled against an “RTE Contract” phase application header file (i.e. object-code SW-Cs) generated in compatibility mode shall be compatible with an RTE that was generated in compatibility mode.

The use of well-defined data structures imposes tight constraints on the RTE implementation and therefore restricts the freedom of RTE vendors to optimize the solution of object-code components but have the advantage that RTE generators from different vendors can be used to compile a binary-component and to generate the RTE.

Note that even when an RTE generator is operating in compatibility mode the data structures used for *source-code* components are not defined thus permitting vendor-specific optimizations to be applied.

5.1.2.1.2 Vendor Mode

Vendor mode is an optional operating mode where the data structures defined in the “RTE Contract” phase and used in the “RTE Generation” phase are implementation specific rather than “standardized”.

[rte_sws_1152] An RTE generator may optionally support *vendor mode*.

The data structures defined and declared when an RTE generator operates in vendor mode are implementation specific and therefore *not* described in this document. This omission is deliberate and permits vendor-specific optimizations to be implemented for object-code components. It also means that RTE generators from different vendors are unlikely to be compatible when run in the vendor mode.

[rte_sws_1234] An AUTOSAR software-component shall be assumed to be operating in “compatibility” mode unless “vendor mode” is explicitly requested.

The potential for more efficient implementations of object-code components offered by the vendor mode comes at the expense of requiring high cohesion between object-code components (compiled after the “RTE Contract” phase) and the generated RTE. However, this is not as restrictive as it may seem at first sight since the tight coupling is also reflected in many other aspects of the AUTOSAR methodology, not least of which is the requirement that the same compiler (and compatible options) is used when compiling both the object-code component and the RTE.

5.1.2.2 Optimization Modes

The actual RTE code is generated – based on the input information – for each ECU individually. To allow optimization during the RTE generation one of the two general optimization directions can be specified: **MEMORY** consumption or execution **RUNTIME**.

[rte_sws_5053] The RTE Generator shall optimize the generated RTE code either for memory consumption or execution runtime depending on the provided input information (see `rte_sws_in_5060`).

5.1.3 Generator external configuration switches

There are use-cases where there is need to influence the behavior of the RTE Generator without changing the RTE Configuration description. In order to support such use-cases this section collects the *external configuration switches*.

Note: it is not specified how these switches shall be implemented in the actual RTE Generator implementation.

[rte_sws_5099] The RTE Generator shall support the *external configuration switch* `strictUnconnectedRPortCheck` which, when enabled, forces the RTE Generator to consider unconnected R-Ports as an error.

5.2 API Principles

[rte_sws_1316] The RTE shall be configured and/or generated for each ECU [RTE00021].

Part of the process is the customization (i.e. configuration or generation) of the RTE API for each AUTOSAR software-component on the ECU. The customization of the API implementation for each AUTOSAR software-component, whether by generation anew or configuration of library code, permits improved run-time efficiency and reduces memory overheads.

The design of the RTE API has been guided by the following core principles:

- The API should be orthogonal – there should be only one way of performing a task.
- **[rte_sws_1314]** The API shall be compiler independent.
- **[rte_sws_3787]** The RTE implementation shall use the compiler abstraction.
- **[rte_sws_1315]** The API shall support components where the source-code is available [RTE00024] and where only object-code is available [RTE00140].
- The API shall support the multiple instantiation of AUTOSAR software-components [RTE00011] that share code [RTE00012].

Two forms of the RTE API are available to software-components; direct and indirect. The direct API has been designed with regard to efficient invocation and includes an API mapping that can be used by an RTE generator to optimize a component's API, for example, to permit the direct invocation of the generated API functions or even eliding the generated RTE completely. The indirect API cannot be optimized using the API mapping but has the advantage that the handle used to access the API can be stored in memory and accessed, via an iterator, to apply the same API to multiple ports.

5.2.1 RTE Namespace

All RTE symbols (e.g. function names, global variables, etc.) visible within the global namespace are required to use the “Rte” prefix.

[rte_sws_1171] All externally visible symbols created by the RTE generator shall use the prefix `Rte_`.

In order to maintain control over the RTE namespace the creation of symbols in the global namespace using the prefix `Rte_` is reserved for the RTE generator.

The generated RTE is required to work with components written in several source languages and therefore should not use language specific features, such as C++ namespaces, to ensure symbol name uniqueness.

5.2.2 Direct API

The direct invocation form is the form used to present the RTE API in Section 5.6. The RTE direct API mapping is designed to be optimizable so that the instance handle is elided (and therefore imposes zero run-time overhead) when the RTE generator can determine that exactly one instance of a component is mapped to an ECU.

[rte_sws_1048] The RTE shall support direct invocation of generated API functions where the instance handle is passed to the API as the first formal parameter.

All runnable entities for a AUTOSAR software-component type are passed the same instance handle type (as the first formal parameter) and can therefore use the same type definition from the component’s application header file.

The direct API can also be further optimized for source code components via the API mapping.

The direct API is typically implemented as macros that are modified by the RTE generator depending on configuration. This technique places certain restrictions on how the API can be used within a program, for example, it is not possible in C to take the address of a macro and therefore direct API functions cannot be placed within a function table or array. If it is required by the implementation of a software-component to derive a pointer to an object for the port API, for instance to setup a constant function pointer table, the *PortAPIOption enableTakeAddress* can be used. Additionally the indirect API provides support for API addresses and iteration over ports.

[rte_sws_7100] If a *PortPrototype* is referenced by *PortAPIOption* with *enableTakeAddress = TRUE* the RTE generator has to provide “C” functions and non function like macro for the API related to this port.

The *PortAPIOption* attribute *enableTakeAddress = TRUE* is not supported for software-components supporting multiple instantiation. See see `rte_sws_in_7101`.

5.2.3 Indirect API

The indirect API is an optional form of API invocation that uses indirection through a port handle to invoke RTE API functions rather than direct invocation. This form is less efficient (the indirection cannot be optimized away) but supports a different programming style that may be more convenient. For example, when using the indirect API, an array of port handles of the same interface and provide/require direction is provided by RTE and the same RTE API can be invoked for multiple ports by iterating over the array.

Both direct and indirect forms of API call are equivalent and result in the same generated RTE function being invoked.

Whether the indirect API is generated or not can be specified for each software component and for each port prototype of the software component separately with the **indirectAPI** attribute, see `rte_sws.in_3798`.

The semantics of the port handle must be the same in both the “RTE Contract” and “RTE Generation” phases since the port handle accesses the standardized data structures of the RTE.

It is possible to mix the indirect and direct APIs within the same SW-C, if the indirect API is present for the SW-C.

The indirect API uses port handles during the invocation of RTE API calls. The type of the port handle is determined by the port interface that types the port which means that if a component declares multiple ports typed by the same port interface the port handle points to an array of port data structures and the same API invoked for each element.

The port handle type is defined in Section 5.4.2.5.

5.2.3.1 Accessing Port Handles

An AUTOSAR SW-C needs to obtain port handles using the instance handle before the indirect API can be used. The definition of the instance handle in Section 5.4.2 defines the “Port API” section of the component data structure and these entries can be used to access the port handles in either object-code or source-code components.

The API `Rte_Ports` and `Rte_NPorts` provides port data handles of a given interface. Example 5.1 shows how the indirect API can be used to apply the same operation to multiple ports in a component within a loop.

Example 5.1

The port handle points to an array that can be used within a loop to apply the same operation to each port. The following example sends the same data to each receiver:

```
1 void TT1(Rte_Instance self)
2 {
3     Rte_PortHandle_interface1_P my_array;
4     my_array=Rte_Ports_interface1_P(self);
5     int s;
6     for(s = 0; s < Rte_NPorts_interface1_P(self); s++) {
7         my_array[s].Send_a(23);
8     }
9 }
```

Note that if `csInterface1` is a client/server interface with an operation `op`, the mechanism sketched in Example 5.1 only works if `op` is invoked either by all clients synchronously or by all clients asynchronously, since the signature of `Rte_Call` and the existence of `Rte_Result` depend on the kind of invocation (see restriction `rte_sws_3605`).

5.2.4 DataReadAccess and DataWriteAccess

The RTE is required to support `DataReadAccess` and `DataWriteAccess` semantics for data elements. The required semantics are subject to two constraints:

- For `DataReadAccess`, the data accessed by a runnable entity must not change during the lifetime of the runnable entity.
- For `DataWriteAccess`, the data written by a runnable entity is only visible to other runnable entities after the accessing runnable entity has terminated.

The generated RTE satisfies both requirements through data copies that are created when the RTE is generated based on the known task and runnable mapping.

Example 5.2

Consider a data element, `a`, of port `p` which is accessed using `DataReadAccess` semantics by runnable `re1` and `DataWriteAccess` by runnable `re2`. Furthermore, consider that `re1` and `re2` are mapped to different tasks and that execution of `re1` can pre-empt `re2`.

In this example, the RTE will create two different copies to contain `a` to prevent updates from `re2` 'corrupting' the value access by `re1` since the latter must remain unchanged during the lifetime of `re1`.

The RTE API includes three API calls to support `DataReadAccess` and `DataWriteAccess` for a software-component; `Rte_IRead` (see Section 5.6.15), `Rte_IWrite`, and `Rte_IWriteRef` (see Section 5.6.16 and 5.6.17). The API calls `Rte_IRead` and `Rte_IWrite` access the data copies (for read and write access respectively). The API call `Rte_IWriteRef` returns a reference to the data copy, thus enabling the runnable to write the data directly. This is especially useful for complex data types and strings. The use of an API call for reading and writing enables the definition to be changed based on the task and runnable mapping without affecting the software-component code.

Example 5.3

Consider a data element, *a*, of port *p* which is declared as being accessed using `DataWriteAccess` semantics by runnables *re1* and *re2* within component *c*. The RTE API for component *c* will then contain four API functions to write the data element;

```
1 void Rte_IWrite_re1_p_a(Rte_Instance self, <type> val);
2 void Rte_IWrite_re2_p_a(Rte_Instance self, <type> val);
3 <type> Rte_IWriteRef_re1_p_a(Rte_Instance self);
4 <type> Rte_IWriteRef_re2_p_a(Rte_Instance self);
```

The API calls are used by *re1* and *re2* as required. The definitions of the API depend on where the data copies are defined. If both *re1* and *re2* are mapped to the same task then each can access the same copy. However, if *re1* and *re2* are mapped to different (pre-emptable) tasks then the RTE will ensure that each API access a different copy.

The `Rte_IRead` and `Rte_IWrite` use the “data handles” defined in the component data structure (see Section 5.4.2).

5.2.5 PerInstanceMemory

The RTE is required to support `PerInstanceMemory` [RTE00013].

The component’s instance handle defines a particular instance of a component and is therefore used when accessing the `PerInstanceMemory` using the `Rte_Pim` API.

The `Rte_Pim` API does not impose the RTE to apply a data consistency mechanism for the access to `PerInstanceMemory`. An application is responsible for consistency of accessed data by itself. This design decision permits efficient (zero overhead) access when required. If a component possesses multiple runnable entities that require concurrent access to the same `PerInstanceMemory`, an exclusive area can be used to ensure data consistency, either through explicit `Rte_Enter` and `Rte_Exit` API calls or by declaring that, implicitly, the runnable entities run inside an exclusive area.

Thus, the `PerInstanceMemory` is exclusively used by a particular software-component instance and needs to be declared and allocated (statically).

[rte_sws_2303] The generated RTE shall declare `PerInstanceMemory` in accordance to the attribute *type* of a particular *PerInstanceMemory*.

In addition, the attribute *type* needs to be defined in the corresponding software-component header. Therefore, the attribute *typeDefinition* of the *PerInstanceMemory* contains its definition as plain text string. It is assumed that this text is valid ‘C’ syntax, because it will be included verbatim in the application header file.

[rte_sws_2304] The generated RTE shall define the type of a `PerInstanceMemory` by interpreting the text string of the attribute *typeDefinition* of a particular *PerInstanceMemory* as the ‘C’ definition.

Note that the type is specified within the scope of a software component and therefore not necessarily unique within the scope of the ECU. Therefore the RTE needs to define a unique type within the RTE Types header file while providing the component-specific type via the application header file to the software component.

[rte_sws_3789] If there is no AUTOSAR data type in the input configuration with the same name as the value of the `type` attribute of a `PerInstanceMemory`, the RTE types header file shall contain the type definition

```
typedef <typedefinition> Rte_PimType_<c>_<t>;
```

where

- `<typedefinition>` is the value of the `typeDefinition` attribute of the `PerInstanceMemory`,
- `<c>` is the name of the component type to which the `PerInstanceMemory` belongs and
- `<t>` is the value of the `type` attribute of the `PerInstanceMemory`.

[rte_sws_3782] If there is no AUTOSAR data type in the input configuration with the same name as the value of the `type` attribute of a `PerInstanceMemory`, the RTE generator shall define the type in the application header file as

```
typedef Rte_PimType_<c>_<t> <t>;
```

where

- `<c>` is the name of the component type to which the `PerInstanceMemory` belongs and
- `<t>` is the value of the `type` attribute of the `PerInstanceMemory`.

Note that it shall be possible to reuse an AUTOSAR data type for the definition of a `PerInstanceMemory`.

[rte_sws_5195] If there is an AUTOSAR data type in the input configuration with the same name as the value of the `type` attribute of the `PerInstanceMemory`, the RTE Generator shall issue a warning that the data type has possibly been misused.

[rte_sws_2305] The generated RTE shall instantiate (or allocate) declared `PerInstanceMemory`.

[rte_sws_5062] In case the `PerInstanceMemory` is used as a permanent ram mirror for the `NvRam` manager the name for the instantiated `PerInstanceMemory` shall be taken from the input information `RamBlockLocationSymbol` `rte_sws_in_5061`. Otherwise the RTE generator is free to choose an arbitrary name.

Note that the memory allocated for a `PerInstanceMemory` is not initialized by the generated RTE, but by the corresponding software-component instances.

Example 5.4

A software-component `c` contains a particular `PerInstanceMemory` `mem` with the attributes `type = MyMemType` and `typeDefinition =`

```
struct {uint16
val1; uint8 * val2; };. This description shall result in the following
code:
```

In the RTE Types header file:

```
1 /* typedef to ensure unique typename */
2 /* according to the attributes */
3 /* 'type' and 'typeDefinition' */
4 typedef struct{
5     uint16 val1;
6     uint8 * val2;
7 } Rte_PimType_c_MyMemType;
```

In the respective application header file:

```
1 /* typedef visible within the scope */
2 /* of the component according to the attributes */
3 /* 'type' and 'typeDefinition' */
4 typedef Rte_PimType_c_MyMemType MyMemType;
```

In Rte.c:

```
1 /* declare and instantiate mem1 */
2 /* "mem1" name may be taken from RamBlockLocationSymbol */
3 Rte_PimType_c_MyMemType mem1;
```

Note that the name used for the definition of the `PerInstanceMemory` may be used outside of the RTE. One use-case is to support the definition of the link between the NvRam Manager's permanent blocks and the SW-Components. The name in `RamBlockLocationSymbol` (`rte_sws_in_5061`) is used to configure the location at which the NvRam Manager shall store and retrieve the permanent block content. For a detailed description please refer to the SW-Component Template [17].

5.2.6 API Mapping

The RTE API is implemented by macros and generated API functions that are created (or configured, depending on the implementation) by the RTE generator during the "RTE Generation" phase. Typically one customized macro or function is created for each "end" of a communication though the RTE generator may elide or combine custom functions to improve run-time efficiency or memory overheads.

[rte_sws_1274] The API mapping shall be implemented in the application header file.

The RTE generator is required to provide a mapping from the RTE API name to the generated function [RTE00051]. The API mapping provides a level of indirection necessary to support binary components and multiple component instances. The indirection is necessary for two reasons. Firstly, some information may not be known when the component is created, for example, the component's instance name, but are necessary to ensure that the names of the generated functions are unique. Secondly, the names

of the generated API functions should be unique (so that the ECU image can link correctly) and the steps taken to ensure this may make the names not “user-friendly”. Therefore, the primary rationale for the API mapping is to provide the required abstraction that means that a component does not need to concern itself with the preceding problems.

The requirements on the API mapping depend on the phase in which an RTE generator is operating. The requirements on the API mapping are only binding for RTE generators operating in compatibility mode.

5.2.6.1 “RTE Contract” Phase

Within the “RTE Contract” phase the API mapping is required to convert from the source API call (as defined in Section 5.6) to the runnable entity provided by a software-component or the implementation of the API function created by the RTE generator.

When compiled against a “RTE Contract” phase header file a software-component that can be multiply instantiated is required to use a general API mapping that uses the instance handle to access the function table defined in the component data structure.

[rte_sws_3706] If a software-component supports multiple instantiation `rte_sws_in_0004`, the “RTE Contract” phase API mapping shall access the generated RTE functions using the instance handle to indirect through the generated function table in the component data structure.

Example 5.5

For a required client-server port ‘p1’ with operation ‘a’ with a single argument, the general form of the API mapping would be:

```
1 #define Rte_Call_p1_a(s,v) ((s)->p1.Call_a(v))
```

Where `s` is the instance handle.

[rte_sws_3707] If a software-component does not support multiple instantiation `rte_sws_in_0004`, the “RTE Contract” phase API mapping shall access the generated RTE functions directly.

When accessed directly, the names of the generated functions are formed according to the following rules:

- **[rte_sws_1143]** The function generated for API calls of the form `<name>_<p>_<o>` shall be `<name>_<c>_<p>_<o>` where `<name>` is the API root (e.g. Call), `<p>` the port name, `<o>` the data element or operation name and `<c>` the component type name.
- **[rte_sws_1348]** The function generated for API calls of the form `<name>_<re>_<p>_<o>` shall be `<name>_<c>_<re>_<p>_<o>` where `<name>` is the API root (e.g. IrvRead), `<p>` the port name, `<re>` the runnable en-

tity name and `<o>` the data element or operation name and `<c>` the component type name.

- **[rte_sws_1155]** The function generated for API calls of the form `<name>_<e>` shall be `<name>_<c>_<e>` where `<name>` is the API root (e.g. Enter), `<e>` the API name (e.g. an exclusive area name) and `<c>` is the component type name.
- **[rte_sws_1156]** The macro generated for the `Rte_Pim` and `Rte_CData` API calls shall map to the relevant fields of the component data structure.

The functions generated that are the destination of the API mapping, which is created during the “RTE Contract” phase, are created by the RTE generator during the second “RTE Generation” phase.

[rte_sws_1153] The generated function (or runnable) shall take the same parameters, in the same order, as the API mapping.

Example 5.6

For a required client-server port ‘p1’ with operation ‘a’ with a single argument for component type ‘c1’ for which multiple instantiation is forbidden, the following mapping would be generated:

```
1 #define Rte_Call_p1_a Rte_Call_c1_p1_a
```

5.2.6.2 “RTE Generation” Phase

There are no requirements on the *form* that the API mapping created during the “RTE Generation” phase should take. This is because the application header files defined during this phase are used by source-code components and therefore compatibility between the generated RTE and source-code components is automatic.

The RTE generator is required to produce the component data structure instances required by object-code components and multiply instantiated source-code components.

If multiple instantiation of a software-component is forbidden, then the API mapping specified for the “RTE Contract” phase (Section 5.2.6.1) defines the names of the generated functions. If multiple instantiation is possible, there are no corresponding requirements that define the name of the generated function since all accesses to the generated functions are performed via the component data structure which contains well-defined entries (Sections 5.4.2.5 and 5.4.2.5).

5.2.6.3 Function Elidation

Using the “RTE Generation” phase API mapping, it is possible for the RTE generator to elide the use of generated RTE functions.

[rte_sws_1146] If the API mapping elides an RTE function the “RTE Generation” phase API mapping mechanism shall ensure that the invoking component still receives a “return value” so that no changes to the AUTOSAR software-component are necessary.

In C, the elidation of API calls can be achieved using a comma expression¹

Example 5.7

As an example, consider the following component code:

```
1 Std_ReturnType s;  
2 s = Rte_Send_p1_a(self, 23);
```

Furthermore, assume that the communication attributes are specified such that the sender-receiver communication can be performed as a direct assignment and therefore no RTE API call needs to be generated. However, the component source cannot be modified and expects to receive an `Std_ReturnType` as the return. The “RTE Generation” phase API mapping could then be rewritten as:

```
1 #define Rte_Send_p1_a(s,a) (<var> = (a), RTE_E_OK)
```

Where `<var>` is the implementation dependent name for an RTE created cache between sender and receiver.

5.2.6.4 API Naming Conventions

An AUTOSAR software-component communicates with other components (including basic software) through ports and therefore the names that constitute the RTE API are formed from the combination of the API call’s functionality (e.g. Call, Send) that defines the API root name and the access point through which the API operates.

For any API that operates through a port, the API’s access point includes the port name.

A `SenderReceiverInterface` can support multiple data items and a `ClientServerInterface` can support multiple operations, any of which can be invoked through the requiring port by a client. The RTE API therefore needs a mechanism to indicate which data item/operation on the port to access and this is implemented by including the data item/operation name in the API’s access point.

As described above, the RTE API mapping is responsible for mapping the RTE API name to the correct generated RTE function. The API mapping permits an RTE generator to include targeted optimization as well as removing the need to implement functions that act as routing functions from generic API calls to particular functions within the generated RTE.

¹This is contrary to MISRA Rule 42 “*comma expression shall not be used except in the control expression of a for loop*”. However, a comma expression is valid, legal, C and the elidation cannot be achieved without a comma expression and therefore the rule must be relaxed.

For C and C++ the RTE API names introduce symbols into global scope and therefore the names are required to be prefixed with `Rte_rte_sws_1171`.

5.2.6.5 API Parameters

All API parameters fall into one of two classes; parameters that are strictly read-only (“In” parameters) and parameters whose value may be modified by the API function (“In/Out” and “Out” parameters).

The type of these parameters is taken from the data element prototype or operation prototype in the interface that characterizes the port for which the API is being generated.

Some RTE API calls (`Rte_CData`, `Rte_Calprm`, `Rte_IRead`, and `Rte_IWriteRef`) pass the communication data as the return value. In these cases the type of the return value is taken from the `calprm` element prototype or the data element prototype in the interface that characterizes the port for which the API is being generated.

5.2.6.5.1 Primitive Data Types except Strings

- “In” Parameters

[rte_sws_1017] All input parameters that are a primitive data type (with the exception of a string) shall be passed by value.

- “Out” Parameters

[rte_sws_1019] All output parameters that are a primitive data type (with the exception of a string) shall be passed by reference.

- “In/Out” Parameters

[rte_sws_1020] All bi-directional parameters (i.e. both input and output) that are a primitive data type (with the exception of a string) shall be passed by reference.

- “Return” Value

[rte_sws_5195] If communication data of a primitive data type (with the exception of a string) is passed as the return value, the type of the return value shall be the data type specified at the `calprm` element prototype or the data element prototype.

5.2.6.5.2 Record Composite Data Types

- “In” Parameters

[rte_sws_1018] All input parameters that are a record composite data type shall be passed by reference.

- **“Out” Parameters**

[rte_sws_5196] All output parameters that are a record composite data type shall be passed by reference.

- **“In/Out” Parameters**

[rte_sws_5197] All bi-directional parameters (i.e. both input and output) that are a record composite data type shall be passed by reference.

- **“Return” Value**

[rte_sws_5198] If communication data of a record composite data type is passed as the return value, the type of the return value shall be a pointer to the data type specified at the calprm element prototype or the data element prototype.

5.2.6.5.3 Array Composite Data Types and Strings

In previous revisions of this document there existed some confusion and incompatibilities regarding the passing of arrays in the RTE API. Some implementations typed array parameters as a pointer to the array base type, others typed array parameters as a pointer to the array type itself. On object code level both variants are equivalent. On source code level compiling a software component against an Application Header File using the other variant of array passing leads to type incompatibility warnings and error messages.

In order to support the coexistence of both array passing variants, a C preprocessor macro is introduced to select the array passing scheme per software component type:

[rte_sws_5199] The Application Header File shall contain the RTE API macro and function prototype definitions for both array passing schemes.

[rte_sws_5200] The array passing scheme used in the Application Header File the implementation of the software component is compiled against shall be selectable via the C preprocessor define `RTE_PTR2ARRAYBASETYPE_PASSING`.

Dependent on the direction and the passing scheme arrays shall be passed as follows:

- **“In” Parameters**

[rte_sws_5107] If the C preprocessor macro `RTE_PTR2ARRAYBASETYPE_PASSING` is defined, input parameters that are an array composite data type or a string shall be passed as an array expression (that is a pointer to the array base type). This requirement applies to the Application Header File only.

Note that AUTOSAR defines a string as a primitive data type yet due to its inherent size it would be inefficient to pass by value and is therefore treated the same as an array composite data type.

[rte_sws_5201] If the C preprocessor macro `RTE_PTR2ARRAYBASETYPE_PASSING` is undefined, input parameters that are an array composite data type or a string shall be passed as a pointer to the array type. This requirement applies to the Application Header File Only.

[rte_sws_5202] For the RTE implementation files (excluding the Application Header Files) an array passing scheme according either to `rte_sws_5107` or to `rte_sws_5201` shall be implemented. A dependency on the C preprocessor macro `RTE_PTR2ARRAYBASETYPE_PASSING` is not necessary.

Note that for the implementation of the RTE itself it is not necessary to know which of the two array passing variants was selected for the implementation of the software component, because both variants are object-code compatible.

- **“Out” Parameters**

[rte_sws_5108] If the C preprocessor macro `RTE_PTR2ARRAYBASETYPE_PASSING` is defined, output parameters that are an array composite data type or a string shall be passed as an array expression (that is a pointer to the array base type). This requirement applies to the Application Header File Only.

[rte_sws_5203] If the C preprocessor macro `RTE_PTR2ARRAYBASETYPE_PASSING` is undefined, output parameters that are an array composite data type or a string shall be passed as a pointer to the array type. This requirement applies to the Application Header File Only.

[rte_sws_5204] For the RTE implementation files (excluding the Application Header Files) an array passing scheme according either to `rte_sws_5108` or to `rte_sws_5203` shall be implemented. A dependency on the C preprocessor macro `RTE_PTR2ARRAYBASETYPE_PASSING` is not necessary.

- **“In/Out” Parameters**

[rte_sws_5109] If the C preprocessor macro `RTE_PTR2ARRAYBASETYPE_PASSING` is defined, bi-directional parameters (i.e. both input and output) that are an array composite data type or a string shall be passed as an array expression (that is a pointer to the array base type). This requirement applies to the Application Header File Only.

[rte_sws_5205] If the C preprocessor macro `RTE_PTR2ARRAYBASETYPE_PASSING` is undefined, bi-directional parameters (i.e. both input and output) that are an array composite data type or a string shall be passed as a pointer to the array type. This requirement applies to the Application Header File Only.

[rte_sws_5206] For the RTE implementation files (excluding the Application Header Files) an array passing scheme according either to `rte_sws_5109` or to `rte_sws_5205` shall be implemented. A dependency on the C preprocessor macro `RTE_PTR2ARRAYBASETYPE_PASSING` is not necessary.

- **“Return” Value**

[rte_sws_5207] If communication data of an array composite data type is passed as the return value, the type of the return value shall be a pointer to the array base type under the condition, that the C preprocessor macro `RTE_PTR2ARRAYBASETYPE_PASSING` is defined. This requirement applies to the Application Header File Only.

[rte_sws_5208] If communication data of an array composite data type is passed as the return value, the type of the return value shall be a pointer to the array type under the condition, that the C preprocessor macro `RTE_PTR2ARRAYBASETYPE_PASSING` is undefined. This requirement applies to the Application Header File Only.

[rte_sws_5209] For the RTE implementation files (excluding the Application Header Files) an array passing scheme according either to `rte_sws_5207` or to `rte_sws_5208` shall be implemented. A dependency on the C preprocessor macro `RTE_PTR2ARRAYBASETYPE_PASSING` is not necessary.

Note that in case of a n -dimensional array the array base type is the type of an individual array element and not the type of the array of the dimension $n - 1$. E. g. for a type `MyInt32_Array_8_9` the pointer to the array base type is `MyInt32 *` and not `MyInt32_Array_9 *`.

Example 5.8

Consider an RTE API call taking an array as an “out” parameter for a singly instantiated SW-C. The signature of the API in the Application Header File will be:

```

1  #ifdef RTE_PTR2ARRAYBASETYPE_PASSING
2      Std_ReturnType Rte_Write_p1_d1(UInt32* value);
3  #else
4      Std_ReturnType Rte_Write_p1_d1(UInt32Array_8* value);
5  #endif

```

The implementation of a software component passing arrays as pointers to the array base type could invoke the API function as follows:

```

1  #define RTE_PTR2ARRAYBASETYPE_PASSING
2  #include "Rte_swc.h"
3  ...
4  UInt32Array_8 myArray; /* or: UInt32 myArray[8]; */
5  ...
6  Rte_Write_p1_d1(myArray);

```

The implementation of a software component passing arrays as pointers to the array type could invoke the API function as follows:

```

1  #include "Rte_swc.h"
2  ...
3  UInt32Array_8 myArray;
4  ...
5  Rte_Write_p1_d1(&myArray);

```

Example 5.9

Consider an RTE API call returning an array for a singly instantiated SW-C. The signature of the API in the Application Header File will be:

```

1 #define RTE_PTR2ARRAYBASETYPE_PASSING
2     UInt32* Rte_IRead_re_p1_d1(void);
3 #else
4     UInt32Array_8* Rte_IRead_re_p1_d1(void);
5 #endif

```

The implementation of a software component passing arrays as pointers to the array base type could invoke the API function as follows:

```

1 #define RTE_PTR2ARRAYBASETYPE_PASSING
2 #include "Rte_sw.c.h"
3 ...
4 UInt32 value0;
5 UInt32* myArray;
6 ...
7 myArray = Rte_IRead_re_p1_d1();
8 value0 = myArray[0];

```

The implementation of a software component passing arrays as pointers to the array type could invoke the API function as follows:

```

1 #include "Rte_sw.c.h"
2 ...
3 UInt32 value0;
4 UInt32Array_8 myArray;
5 ...
6 myArray = Rte_IRead_re_p1_d1();
7 value0 = (*myArray)[0];

```

5.2.6.6 Error Handling

In RTE, error and status information is defined with the data type `Std_ReturnType`, see Section 5.5.1.

It is possible to distinguish between infrastructure errors and application errors. Infrastructure errors are caused by a resource failure or an invalid input parameter. Infrastructure errors usually occur in the basic software or hardware along the communication path of a data element. Application errors are reported by a SW-C or by AUTOSAR services. RTE has the capability to treat application errors that are forwarded

- by return value in client server communication or
- by signal invalidation in sender receiver communication with `isQueued` set to `false`.

Errors that are detected during an RTE API call are notified to the caller using the API's return value.

[rte_sws_1034] Error states (including 'no error') shall only be passed as return value of the RTE API to the AUTOSAR SW-C.

Requirement `rte_sws_1034` ensures that, irrespective of whether the API is blocking or non-blocking, the error is collected at the same time the data is made available to the caller thus ensuring that both items are accessed consistently.

Certain RTE API calls operate asynchronously from the underlying communication mechanism. In this case, the return value from the API indicates only errors detected during that API call. Errors detected after the API has terminated are returned using a different mechanism `rte_sws_1111`. RTE also provides an 'implicit' API for direct access to virtually shared memory. This API does not return any errors. The underlying communication is decoupled. Instead, an API is provided to pick up the current status of the corresponding data element.

5.2.6.7 Success Feedback

The RTE supports the notification of results of transmission attempts to an AUTOSAR software-component.

The `Rte_Feedback` API `rte_sws_1083` can be configured to return the transmission result as either a blocking or non-blocking API or via activation of a runnable entity.

5.2.7 Unconnected Ports

[rte_sws_1329] The RTE shall handle both require and provide ports that are not connected.

[rte_sws_5100] The handling of require ports as an error shall be configured using `rte_sws_5099`.

The API calls for unconnected ports are specified to behave as if the port was connected but the remote communication point took no action.

Unconnected require ports are regarded by the RTE generator as an invalid configuration (see `rte_sws_3019`) if the strict handling has been enabled (see `rte_sws_5099`).

5.2.7.1 Data Elements

5.2.7.1.1 Explicit Communication

[rte_sws_1330] A non-blocking `Rte_Read` API for an unconnected require port typed by a `SenderReceiverInterface` shall return `RTE_E_OK` code as if a sender was connected but did not transmit anything.

Requirement `rte_sws_1330` applies to elements with "data" semantics (`isQueued = false`) and therefore "last is best" semantics. This means that the initial value will be returned.

[rte_sws_1331] A blocking `Rte_Receive` API for an unconnected require port typed by a `SenderReceiverInterface` shall return `RTE_E_TIMEOUT` immediately without waiting for expiry of the timeout.

[rte_sws_1336] A non-blocking `Rte_Receive` API for an unconnected require port typed by a `SenderReceiverInterface` shall return `RTE_E_NO_DATA` immediately.

The existence of blocking and non-blocking `Rte_Read` and `Rte_Receive` API calls is controlled by the presence of `DataReceivePoints`, `DataReceiveEvents` and `WaitPoints` within the SW-C description `rte_sws_1288`, `rte_sws_1289` and `rte_sws_1290`.

[rte_sws_1344] A blocking `Rte_Feedback` API for a `DataElementPrototype` of an unconnected provide port shall return `RTE_E_TRANSMIT_ACK` immediately.

[rte_sws_1345] A non-blocking `Rte_Feedback` API for a `DataElementPrototype` of an unconnected provide port shall return `RTE_E_TRANSMIT_ACK` immediately.

The existence of blocking and non-blocking `Rte_Feedback` API is controlled by the presence of `DataSendPoints`, `DataSendCompletedEvents` and `WaitPoints` within the SW-C description for a `DataElementPrototype` with acknowledgement enabled, see `rte_sws_1283`, `rte_sws_1284`, `rte_sws_1285` and `rte_sws_1286`.

[rte_sws_1332] The `Rte_Send` or `Rte_Write` API for an unconnected provide port typed by a `SenderReceiverInterface` shall discard the input parameters and return `RTE_E_OK`.

The existence of `Rte_Send` or `Rte_Write` is controlled by the presence of `DataSendPoints` within the SW/C description `rte_sws_1280` and `rte_sws_1281`.

[rte_sws_3783] The `Rte_Invalidate` API for an unconnected provide port typed by a `SenderReceiverInterface` shall return `RTE_E_OK`.

The existence of `Rte_Invalidate` is controlled by the presence of `DataSendPoints` within the SW/C description for a `DataElementPrototype` which is marked as invalidatable and has `canInvalidate` enabled `rte_sws_1282`.

5.2.7.1.2 Implicit Communication

[rte_sws_1346] An `Rte_IRead` API for an unconnected require port typed by a `SenderReceiverInterface` shall return the initial value.

The existence of `Rte_IRead` is controlled by the presence of `DataReadAccess` in the SW-C description `rte_sws_1301`.

[rte_sws_1347] An `Rte_IWrite` API for an unconnected provide port typed by a `SenderReceiverInterface` shall discard the written data.

The existence of `Rte_IWrite` is controlled by the presence of `DataWriteAccess` in the SW-C description `rte_sws_1302`.

[rte_sws_3784] An `Rte_IInvalidate` API for an unconnected provide port typed by a `SenderReceiverInterface` shall perform no action.

The existence of `Rte_IInvalidate` is controlled by the presence of `DataWriteAccess` in the SW-C description for a `DataElementPrototype` which is marked as invalidatable and has `canInvalidate` enabled `rte_sws_3801`.

[rte_sws_3785] An `Rte_IStatus` API for an unconnected require port typed by a `SenderReceiverInterface` shall return `RTE_E_OK`.

The existence of `Rte_IStatus` is controlled by the presence of `DataReadAccess` in the SW-C description for a `DataElementPrototype` with data element outdated notification or data element invalidation `rte_sws_2600`.

5.2.7.2 Mode Ports

For the mode user an unconnected mode port behaves as if it was connected to a mode manager that never sends a mode switch notification.

[rte_sws_2638] A `Rte_Mode` API for an unconnected mode port of a mode user shall return the initial state.

[rte_sws_2639] Regarding the modes of an unconnected mode port of a mode user, the mode disabling dependencies on the initial mode shall be permanently active and the mode disabling dependencies on all other modes shall be inactive.

[rte_sws_2640] Regarding the modes of an unconnected mode port of a mode user, RTE will only generate a `ModeSwitchEvent` for entering the initial mode which occurs directly after startup.

[rte_sws_2641] The `Rte_Switch` API for an unconnected mode port of the mode manager shall discard the input parameters and return `RTE_E_OK`.

[rte_sws_2642] A blocking or non blocking `Rte_Feedback` API for an unconnected mode port of the mode manager shall return `RTE_E_OK` immediately.

5.2.7.3 Client-Server

[rte_sws_1333] The `Rte_Result` API for an unconnected asynchronous require port typed by a `ClientServerInterface` with a `WaitPoint` for the `AsynchronousServerCallReturnsEvent` shall return `RTE_E_TIMEOUT` immediately without waiting for expiry of the timeout.

[rte_sws_1337] The `Rte_Result` API for an unconnected asynchronous require port typed by a `ClientServerInterface` without a `WaitPoint` for the `AsynchronousServerCallReturnsEvent` shall return `RTE_E_NO_DATA` immediately.

[rte_sws_1334] An asynchronous `Rte_Call` API for an unconnected require port typed by a `ClientServerInterface` shall return `RTE_E_OK` immediately.

5.2.8 Non-identical ports

Two ports are permitted to be connected provided that they are characterized by compatible, but not necessarily identical, interfaces. For the full definition of whether two interfaces are compatible, see the System Template.

[rte_sws_1368] The RTE generator must report an error if two connected ports are connected by incompatible interfaces.

A significant issue in determining whether two interfaces are compatible is that the interface characterizing the require port may be a strict subset of the interface characterizing the provide port. This means that there may be provided data elements or operations for which there is no corresponding element in the require port. This can be imagined as a multi-strand wire between the two ports (the assembly connector) where each strand represents the connection between two data elements or operations, and where some of the strands from the 'provide' end are not connected to anything at the 'require' end.

Define, for the purposes of this section, an "unconnected element" as a data element or operation that occurs in the provide interface, but for which no corresponding data element or operation occurs in a particular R-Port's interface.

[rte_sws_1369] For each data element or operation within the provide interface, every connected requirer with an "unconnected element" must be treated as if it were not connected.

Note that requirement `rte_sws_1369` means that in the case of a 1:n Sender-Receiver the `Rte_Write` call may transmit to some but not all receivers. Similarly, there may be some clients that cannot write into a server's queue.

The extreme is if all connected requirers have an "unconnected element":

[rte_sws_1370] For a data element or operation in a provide interface which is an unconnected element in every connected R-Port, the generated Send or Write API must act as if the port were unconnected."

See Section 5.2.7 for the required behaviour in this case.

5.3 RTE Modules

Figure 5.1 defines the relationship between header files and how those files are included by modules implementing AUTOSAR software-components and by general, non-component, code.

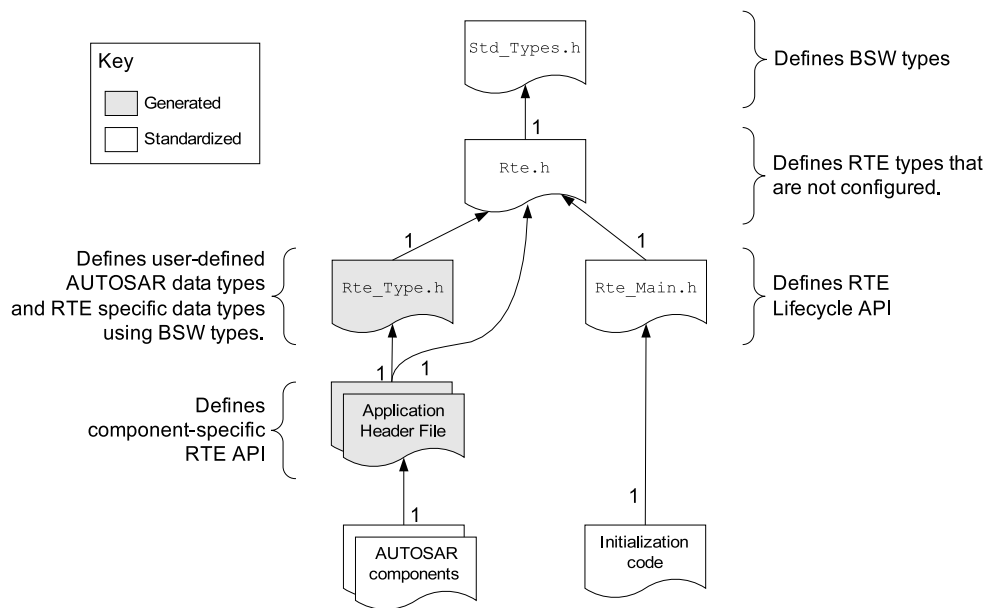


Figure 5.1: Relationships between RTE Header Files

The output of an RTE generator can consist of both generated code and configuration for “library” code that may be supplied as either object code or source code. Both configured and generated code reference standard definitions that are defined in one of two standardized header files; the *RTE Header File* and the *Lifecycle Header File*.

The relationship between the RTE header file, application header files, the lifecycle header file and AUTOSAR software-components is illustrated in Figure 5.1.

5.3.1 RTE Header File

The RTE header file defines fixed elements of the RTE that do not need to be generated or configured for each ECU.

[rte_sws_1157] For C/C++ AUTOSAR software-components, the name of the RTE header file shall be `Rte.h`.

Typically the contents of the standardized header file are fixed for any particular implementation and therefore it is not created by the RTE generator. However, customization for each generated RTE is not forbidden.

[rte_sws_1164] The RTE header file shall include the file `Std_Types.h`.

The file `Std_Types.h` is the standard AUTOSAR file [12] that defines basic data types including platform specific definitions of unsigned and signed integers and provides access to the compiler abstraction.

The contents of the RTE header file are not restricted to standardized elements that are defined within this document – it can also contain definitions specific to a particular implementation.

5.3.2 Lifecycle Header File

The Lifecycle header file defines the two RTE Lifecycle API calls `Rte_Start` and `Rte_Stop` (see Section 5.8).

[rte_sws_1158] For C/C++ AUTOSAR software-components, the name of the lifecycle header file shall be `Rte_Main.h`.

[rte_sws_1159] The lifecycle header file shall include the *RTE header file*.

5.3.3 Application Header File

The application header file [RTE00087] is central to the definition of the RTE API. An application header file defines the RTE API and any associated data structures that are required by the RTE implementation. But the application header file is not allowed to create objects in memory.

[rte_sws_1000] The RTE generator shall create an application header file for each software-component type (excluding `CalprmComponentTypes`) defined in the input.

[rte_sws_3786] The application header file shall not contain code that creates objects in memory.

Due to the restriction `rte_sws_5034` it is only allowed to have exactly one `InternalBehavior` for each component type.

RTE generation consists of two phases; an initial “RTE Contract” phase and a second “RTE Generation” phase (see Section 2.3). Object-code components are compiled after the first phase of RTE generation and therefore the application header file should conform to the form of definitions defined in Sections 5.4.1 and 5.5.2. In contrast, source-code components are compiled after the second phase of RTE generation and therefore the RTE generator produces an optimized application header file based on knowledge of component instantiation and deployment.

5.3.3.1 File Name

[rte_sws_1003] The name of the application header file shall be formed by prefixing the AUTOSAR software-component type name with `Rte_` and appending the result with `.h`.

Example 5.10

The following declaration in the input XML:

```

1 <ATOMIC-SOFTWARE-COMPONENT-TYPE>
2     <SHORT-NAME>Source</SHORT-NAME>
3 </ATOMIC-SOFTWARE-COMPONENT-TYPE>

```

should result in the application header file `Rte_Source.h` being generated.

The component type name is used rather than the component instance name for two reasons; firstly the same component code is used for all component instances and, secondly, the component instance name is an internal identifier, and should not appear outside of generated code.

5.3.3.2 Scope

[rte_sws_1004] The application header file for a component shall contain only information relevant to that component.

Requirement `rte_sws_1004` means that compile time checks ensure that a component that uses the application header file only accesses the generated data structures and functions to which it has been configured. Any other access, e.g. to fields not defined in the customized data structures or RTE API, will fail with a compiler error [RTE00017].

[rte_sws_1005] The application header file shall be valid for both C and C++ source.

Requirement `rte_sws_1005` is met by ensuring that all definitions within the application header file are defined using C linkage if a C++ compiler is used.

[rte_sws_3709] All definitions within in the application header file shall be preceded by the following fragment;

```

1 #ifdef __cplusplus
2 extern "C" {
3 #endif /* __cplusplus */

```

[rte_sws_3710] All definitions within the application header file shall be suffixed by the following fragment;

```

1 #ifdef __cplusplus
2 } /* extern "C" */
3 #endif /* __cplusplus */

```

```

1 #include <Rte_c1.h>
2
3 void
4 runnable_entry(Rte_Instance self)
5 {
6     /* ... server code ... */
7 }

```

Figure 5.2: Skeleton server runnable entity

The definitions of the RTE API contained in the application header file can be optimized during the “RTE Generation” phase when the mapping of software-components to ECUs and the communication matrix is known. Consequently multiple application header files must not be included in the same source module to avoid conflicting definitions of the RTE API definitions that the files contains.

Figure 5.2 illustrates the code structure for the declaration of the entry point of a runnable entity that provides the implementation for a ServerPort in component `c1`. The RTE generator is responsible for creating the API and tasks used to execute the server and the symbol name of the entry point is extracted from the attribute symbol of the runnable entity. The example shows that the first parameter of the entry point function is the software-component’s instance handle `rte_sws_1016`.

Figure 5.2 includes the component-specific application header file `Rte_c1.h` created by the RTE generator. The RTE generator will also create the supporting data structures and the task body to which the runnable is mapped.

The RTE is also responsible for preventing conflicting concurrent accesses when the runnable entity implementing the server operation is triggered as a result of a request from a client received via the communication service or directly via inter-task communication.

5.3.3.3 File Contents

Multiple application header file must not be included in the same module (`rte_sws_1004`) and therefore the file contents should contain a mechanism to enforce this requirement.

[rte_sws_1006] An application header file shall include the following mechanism before any other definitions.

```

1 #ifndef RTE_APPLICATION_HEADER_FILE
2 #error Multiple application header files included.
3 #endif /* RTE_APPLICATION_HEADER_FILE */
4 #define RTE_APPLICATION_HEADER_FILE

```

The RTE uses an instance handle to identify different instances of the same component type. The definition of the instance handle type `rte_sws_1148` is unique to each component type and therefore should be included in the application header file.

[rte_sws_1007] The application header file shall define the type of the instance handle for the component.

All runnable entities for a component are passed the same instance handle type (as the first formal parameter `rte_sws_1016`) and can therefore use the same type definition from the component's application header file.

[rte_sws_1263] The application header file shall include the *AUTOSAR Types Header File*.

The name of the AUTOSAR Types Header File is defined in Section 5.3.4.

The application header file also includes a prototype for each runnable entity entry point (`rte_sws_1132`) and the API mapping (`rte_sws_1274`).

[rte_sws_5078] The application header file shall define the init value of unqueued primitive type `DataElementPrototypes`

```
1 #define Rte_InitValue_<Port>_<DEPType> ((<DataType>) <initValue>)
```

where `<Port>` is the `PortPrototype` `shortName`, `<DEPType>` is the `shortName` of the `DataElementPrototype`, `<DataType>` is the `shortName` of the `DataElementPrototype`'s type and `<initValue>` is the `initValue` specified in the `UnqueuedReceiverComSpec` respectively `UnqueuedSenderComSpec`.

Note that the `initValue` defined may be subject to change due to the fact that for COM configuration it may be possible to change this value during ECU Configuration or even post-build time.

5.3.3.3.1 RTE-Component Interface

The application header file defines the "interface" between a component and the RTE. The interface consists of the RTE API for the component and the prototypes for runnable entities. The definition of the RTE API requires that both relevant data structures and API calls are defined.

The data structures required to support the API are defined in the RTE Types header file `rte_sws_3713`. This enables the definitions to be available to multiple modules to support direct function invocation.

The data structure types are declared in the RTE Types file whereas the instances are defined in the generated RTE. The necessary data structures for object-code software-components are defined 5.5.2.

[rte_sws_1009] The application header file shall define the mapping from the RTE API to the generated API functions that are generated/configured for the component.

The RTE generator is required `rte_sws_1004` to limit the contents of the application header file to only that information that is relevant to that component type. This requirement includes the definition of the API mapping.

[rte_sws_1276] Only RTE API calls that are valid for the particular software-component type shall be defined within the component's application header file.

Requirement `rte_sws_1276` ensures that attempts to invoke invalid API calls will be rejected as a compile-time error [RTE00017].

5.3.4 AUTOSAR Types Header File

The AUTOSAR types header file defines RTE specific types derived either from the input configuration or from the RTE implementation.

The generated RTE can include zero or more AUTOSAR data types created from the definitions of AUTOSAR meta-model classes within the RTE generator's input. The available meta-model classes are defined by the AUTOSAR software-component template and include classes for defining integers, floats as well as "complex" data types such as records.

[rte_sws_1160] The RTE generator shall create the AUTOSAR Types header file defining the AUTOSAR data types and RTE implementation types.

The AUTOSAR data types header file should be output for "RTE Contract" and "RTE Generation" phases. RTE implementation types include the Component Data Structure (Section 5.4.2).

5.3.4.1 File Contents

[rte_sws_2648] The AUTOSAR Types header file shall include the definitions of all AUTOSAR data types irrespective of their use by the generated RTE.

This requirement ensures the availability of AUTOSAR data types for the internal use in AUTOSAR software components.

The types header file may need to define types in terms of BSW types (from the file `Std_Types.h`) or from the implementation specific RTE header file. However, since the RTE header file includes the file `Std_Types.h` already so only the RTE header file needs direct inclusion within the types header file.

[rte_sws_1163] The AUTOSAR Types header file shall include the *RTE header file*.

5.3.4.2 Primitive AUTOSAR Data Types

The AUTOSAR types file defines the mapping from primitive AUTOSAR data-types (defined in the XML) to programming language specific type definitions. The mapping from primitive AUTOSAR data-types to BSW standard types (as defined in `Std_Types.h` is defined in Table 5.1).

Requirement	Meta-type	Range	Base Type
[rte_sws_1175]	CHAR-TYPE	Encoding 'UTF-8'	uint8
[rte_sws_1176]	STRING-TYPE	Declaration, <i>n</i> is defined maximum length including zero terminator	uint8[<i>n</i>]
[rte_sws_1177]	INTEGER-TYPE	[-128,127]	sint8
[rte_sws_1178]	INTEGER-TYPE	[-32768,32767]	sint16
[rte_sws_1179]	INTEGER-TYPE	[-2147483648,2147483647]	sint32
[rte_sws_1180]	INTEGER-TYPE	[0,255]	uint8
[rte_sws_1181]	INTEGER-TYPE	[0,65535]	uint16
[rte_sws_1182]	INTEGER-TYPE	[0,4294967295]	uint32
[rte_sws_1183]	OPAQUE-TYPE	Bit length 1..8	uint8
[rte_sws_1184]	OPAQUE-TYPE	Bit length 9..16	uint16
[rte_sws_1185]	OPAQUE-TYPE	Bit length 17..32	uint32
[rte_sws_1186]	REAL-TYPE	Encoding single	float32
[rte_sws_1187]	REAL-TYPE	Encoding double	float64
[rte_sws_1188]	BOOLEAN-TYPE	N/A	boolean

Table 5.1: C/C++ mapping from primitive AUTOSAR data-types

An integer type is defined using either an *open* or *closed* interval – a closed interval includes its endpoints whereas an open interval does not. For simplicity, Table 5.1 defines mappings for integer types using *closed* intervals.

[rte_sws_1265] Where the range expressed in a type definition is not exactly the same as a range defined in Table 5.1, the RTE generator shall select the smallest suitable base type.

Example 5.11 describes the definition of an 11-bit unsigned integer type in terms of a 16-bit base type.

Example 5.11

The following declaration of the user-defined type `UInt11` in the input XML:

```

1  <INTEGER-TYPE>
2    <SHORT-NAME>UInt11</SHORT-NAME>
3    <LOWER-LIMIT>
4      <INTERVAL-TYPE>CLOSED</INTERVAL-TYPE>
5      <VALUE>0</VALUE>
6    </LOWER-LIMIT>
7    <UPPER-LIMIT>
8      <INTERVAL-TYPE>OPEN</INTERVAL-TYPE>
9      <VALUE>2048</VALUE>
10   </UPPER-LIMIT>
11  </INTEGER-TYPE>
    
```

Should result in a mapping to the base type `uint16` and the following type definition;

```
typedef uint16 UInt11;
```

[rte_sws_1214] An attempt to declare a type with a range which cannot be represented by a base type from Table 5.1 shall be rejected by the RTE generator.

Table 5.1 applies to the standard AUTOSAR types as well as user-defined types and primitive data-types with semantics. Using the requirements defined in Table 5.1 the standard AUTOSAR primitive types are mapped as follows:

AUTOSAR Type	BSW Type
UInt4	uint8
SInt4	sint8
UInt8	uint8
SInt8	sint8
UInt16	uint16
SInt16	sint16
UInt32	uint32
SInt32	sint32
Float_with_NaN	float32
Float	float32
Double_with_NaN	float64
Double	float64
Boolean	boolean
Char8	uint8
Char16	uint16

Table 5.2: C/C++ mapping for standard AUTOSAR data-types

5.3.4.3 Complex AUTOSAR Data Types

In addition to the primitive data-types defined in the previous section, it is also necessary for the RTE generator to define complex data-types; arrays and records.

An array definition needs three pieces of information; the array base type, the array name and the number of elements.

[rte_sws_1189] An `ARRAY-TYPE` data-type shall be declared as `typedef <type> <name>[n]` where `<type>` is the base type, `<name>` the data-type name and `n` the number of elements.

Example 5.12

The array data-type declaration;

```

1 <ARRAY-TYPE>
2   <SHORT-NAME>array</SHORT-NAME>
3   <DESC>array of myInt values</DESC>
4   <ELEMENT-TYPE-REF>myInt</ELEMENT-TYPE-REF>
5   <MAX-NUMBER-OF-ELEMENTS>2</MAX-NUMBER-OF-ELEMENTS>
6 </ARRAY-TYPE>

```

Produces the following type definition;

```

1 typedef myInt array[2];

```

ANSI C does not allow a type declaration to have zero elements and therefore we require that the “number of elements” to be a positive integer.

[rte_sws_1190] The number of elements of an `ARRAY_TYPE` data type shall be an integer that is ≥ 1 .

A record definition contains references to one or more data elements with a base type for each element. A record definition is recursive; a data element can include a type reference that is itself another record definition.

[rte_sws_1191] A `RECORD-TYPE` data-type shall be declared as `typedef struct { <elements> } <name>` where `<elements>` is the record element specification and `<name>` the data-type name.

ANSI C does not allow a `struct` to have zero elements and therefore we require that a record include at least one element.

[rte_sws_1192] A record shall include at least one element.

Example 5.13

The record data-type declaration;

```

1 <RECORD-TYPE>
2   <SHORT-NAME>R2</SHORT-NAME>
3   <ELEMENTS>
4     <RECORD-ELEMENT>
5       <SHORT-NAME>Abc</SHORT-NAME>
6       <TYPE-TREF>myBool</TYPE-TREF>
7     </RECORD-ELEMENT>
8     <RECORD-ELEMENT>
9       <SHORT-NAME>Def</SHORT-NAME>
10      <TYPE-TREF>myInt</TYPE-TREF>
11    </RECORD-ELEMENT>
12  </ELEMENTS>
13 </RECORD-TYPE>

```

Produces the following type definition;

```

1 typedef struct {
2   myBool Abc;

```

```

3     myInt    Def;
4 } R2;

```

5.3.4.4 C/C++

The following requirements apply to RTEs generated for C and C++.

[rte_sws_1161] The name of the AUTOSAR types header file shall be `Rte_Type.h`.

[rte_sws_1162] Within the AUTOSAR types header file, each data type shall be defined using `typedef`.

A `typedef` is used when defining a new data type instead of a `#define` even though C only provides weak type checking since other static analysis tools can then be used to overlay strong type checking onto the C before it is compiled and thus detect type errors before the module is even compiled.

5.3.5 VFB Tracing Header File

The VFB Tracing Header File defines the configured VFB Trace events.

[rte_sws_1319] The VFB Tracing Header File shall be created by the RTE Generator during “RTE Generation” phase only.

The VFB Tracing Header file is included by the generated RTE and by the user in the module(s) that define the configured hook functions. The header file includes prototypes for the configured functions to ensure consistency between the invocation by the RTE and the definition by the user.

5.3.5.1 C/C++

The following requirements apply to RTEs generated for C and C++.

[rte_sws_1250] The name of the VFB Tracing Header File shall be `Rte_Hook.h`.

5.3.5.2 File Contents

[rte_sws_1251] The VFB Tracing header file shall include the *RTE Configuration file* (Section 5.3.6).

[rte_sws_1357] The VFB Tracing header file shall include the *AUTOSAR Types Header file* (Section 5.3.4).

[rte_sws_3607] The VFB Tracing header file shall include `Os.h`.

[rte_sws_1320] The VFB Tracing header file shall contain the following code immediately after the include of the RTE Configuration file.

```
1 #ifndef RTE_VFB_TRACE
2 #define RTE_VFB_TRACE (0)
3 #endif /* RTE_VFB_TRACE */
```

Requirement `rte_sws_1320` enables VFB tracing to be globally enabled/disabled within the RTE Configuration file and ensures that it defaults to 'disabled'.

[rte_sws_1236] For each trace event hook function defined in Section 5.10.2, the RTE generator shall define the following code sequence in the VFB Tracing header file:

```
1 #if defined(<trace event>) && (RTE_VFB_TRACE == 0)
2 #undef <trace event>
3 #endif
4 #if defined(<trace event>)
5 #undef <trace event>
6 extern void <trace event>(<params>);
7 #else
8 #define <trace event>(<params>) ((void) (0))
9 #endif /* <trace event> */
```

In the example above, `<trace event>` is the name of trace event hook function and `<params>` is the list of parameter names of the trace event hook function prototype as defined in Section 5.10.2.

The code fragment within `rte_sws_1236` benefits from a brief analysis of its structure. The first `#if` block ensures that an individually configured trace event in the RTE Configuration file `rte_sws_1324` is disabled if tracing is globally disabled `rte_sws_1323`. The second `#if` block emits the prototype for the hook function only if enabled in the RTE Configuration file and thus ensures that only configured trace events are prototyped. The `#undef` is required to ensure that the trace event function is invoked as a function by the generated RTE. The `#else` block comes into effect if the trace event is disabled, either individually `rte_sws_1325` or globally, and ensures that it has no run-time effect. Within the `#else` block the definition to `((void) (0))` enables the hook function to be used within the API Mapping in a comma-expression.

An individual trace event defined in Section 5.10.2 actually defines a class of hook functions. A member of the class is created for each RTE object created (e.g. for each API function, for each task) and therefore an individual trace event may give rise to many hook function definitions in the VFB Tracing header file.

Example 5.14

Consider an API call `Rte_Write_p1_a` for an instance of SW-C `c`. This will result in two trace event hook functions being created by the RTE generator:

```
1 Rte_WriteHook_c_p1_a_Start
```

and

```
1 Rte_WriteHook_c_pl_a_Return
```

5.3.6 RTE Configuration Header File

The RTE Configuration Header file contains user definitions that affect the behaviour of the generated RTE.

The directory containing the required RTE Configuration header file should be included in the compiler's include path when using the VFB tracing header file.

5.3.6.1 C/C++

The following requirements apply to RTEs generated for C and C++.

[rte_sws_1321] The name of the RTE Configuration Header File shall be `Rte_Cfg.h`.

5.3.6.2 File Contents

[rte_sws_1322] The RTE generator shall globally enable VFB tracing when `RTE_VFB_TRACE` is defined in the RTE configuration header file as a non-zero integer.

Note that, as observed in Section 5.10, VFB tracing enables debugging of software components, not the RTE itself.

[rte_sws_1323] The RTE generator shall globally disable VFB tracing when `RTE_VFB_TRACE` is defined in the RTE configuration header file as 0.

As well as globally enabling or disabling VFB tracing, the RTE Configuration header file also configures those individual VFB tracing events that are *enabled*.

[rte_sws_1324] The RTE generator shall enable VFB tracing for a given hook function when there is a `#define` in the RTE configuration header file for the hook function name and tracing is globally enabled.

Note that the particular value assigned by the `#define`, if any, is not significant.

[rte_sws_1325] The RTE generator shall disable VFB tracing for a given hook function when there is no `#define` in the RTE configuration header file for the hook function name even if tracing is globally enabled.

Example 5.15

Consider the trace events from Example 5.14. The trace event for API start is enabled by the following definition;

```
1 #define Rte_WriteHook_il_pl_a_Start
```

And the trace event for API termination is enabled by the following definition;

```
1 #define Rte_WriteHook_i1_p1_a_Return
```

5.3.7 Generated RTE

Figure 5.1 defines the relationship between generated and standardized header files. It is **not** necessary to standardize the relationship between the C module, `Rte.c`, and the header files since when the RTE is generated the application header files are created anew along with the RTE. This means that details of which header files are included by `Rte.c` can be left as an implementation detail.

5.3.7.1 Header File Usage

[rte_sws_1257] In compatibility mode, the Generated RTE module shall include `Os.h`.

[rte_sws_3794] In compatibility mode, the generated RTE module shall include `Com.h`.

[rte_sws_1279] In compatibility mode, the Generated RTE module shall include `Rte.h`.

[rte_sws_1326] In compatibility mode, the Generated RTE module shall include the VFB Tracing header file.

[rte_sws_3788] The generated RTE shall use the file `MemMap.h`.

Figure 5.3 provides an example of how the RTE header and generated header files could be used by a generated RTE.

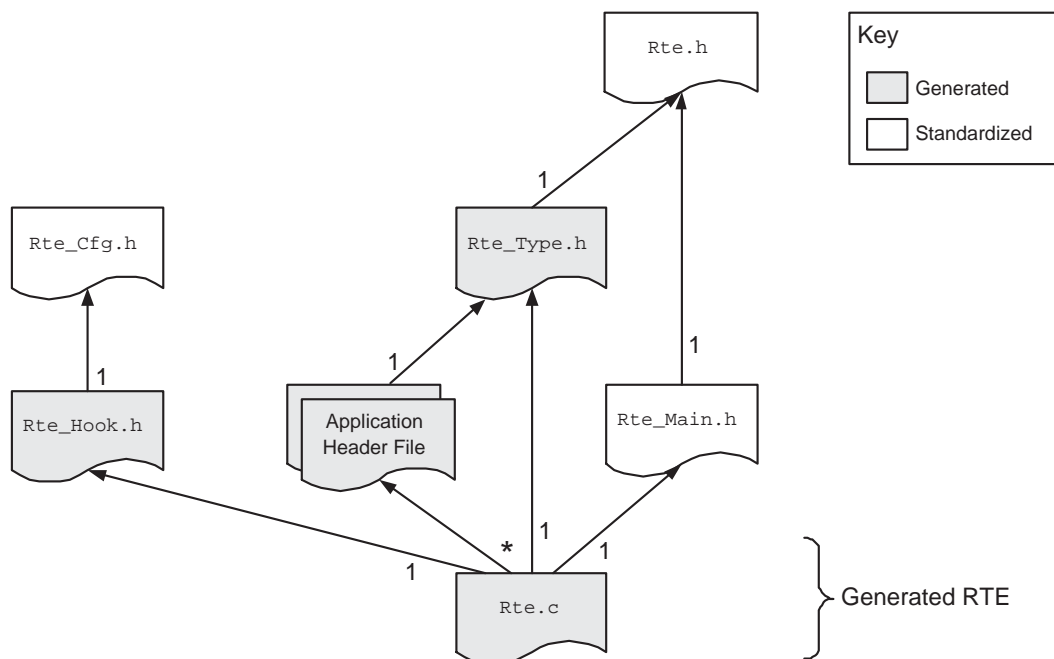


Figure 5.3: Example of header file use by the generated RTE.

In the example in Figure 5.3, the generated RTE C module requires access to the data structures created for each AUTOSAR software-component and therefore includes each application header file². In the example, the generated RTE also includes the RTE header file and the lifecycle header file in order to obtain access to RTE and lifecycle related definitions.

5.3.7.2 C/C++

The following requirements apply to RTEs generated for C and C++.

[rte_sws_1169] The name of the C module containing the generated RTE shall be `Rte.c`.

An RTE that includes configured code from an object-code or source-code library may use additional modules.

5.3.7.3 File Contents

By its very nature the contents of the generated RTE is largely vendor specific. It is therefore only possible to define those common aspects that are visible to the “outside world” such as the names of generated APIs and the definition of component data structures that apply any operating mode.

5.3.7.3.1 Component Data Structures

The *Component Data Structure* (Section 5.4.2) is a per-component data type used to define instance specific information required by the generated RTE.

[rte_sws_3711] The generated RTE shall contain an instance of the relevant Component Data Structure for each software-component instance on the ECU for which the RTE is generated.

[rte_sws_3712] The name of a Component Data Structure instantiated by the RTE generator shall be `Rte_Instance_<name>` where `<name>` is an automatically generated name, created in some manner such that all instance data structure names are unique.

The software component instance name referred to in `rte_sws_3712` is never made visible to the users of the generated RTE. There is therefore no need to specify the precise form that the unique name takes. The `Rte_Instance_` prefix is mandated in order to ensure that no name clashes occur and also to ensure that the structures are readily identifiable in map files, debuggers, etc.

²The requirement that a software module include at most one application header file applies only to modules that actually implement a software-component and therefore does not apply to the generated RTE.

5.3.7.3.2 Generated API

[rte_sws_1266] The RTE module shall define the generated functions that will be invoked when an AUTOSAR software-component makes an RTE API call.

The semantics of the generated functions are not defined (since these will obviously vary depending on the RTE API call that it is implementing) nor are the implementation details (which are vendor specific). However, the names of the generated functions defined in Section 5.2.6.1.

The signature of a generated function is the same as the signature of the relevant RTE API call (see Section 5.6) with the exception that the instance handle can be omitted since the generated function is applicable to a specific software-component instance.

5.3.7.3.3 Callbacks

In addition to the generated functions for the RTE API, the RTE module includes callbacks invoked by COM when signal events (receptions, transmission acknowledgment, etc.) occur.

[rte_sws_1264] The RTE module shall define COM callbacks for relevant signals.

The required callbacks are defined in Section 5.9.2.

[rte_sws_3795] The RTE generator shall generate a separate header file containing the prototypes of the COM callback functions.

[rte_sws_3796] The name of the header file containing the COM callback prototypes shall be `Rte_Cbk.h` in a C/C++ environment.

5.3.7.3.4 Task bodies

The RTE module define task bodies for tasks created by the RTE generator only in compatibility mode.

[rte_sws_1277] In compatibility mode `rte_sws_1257`, the RTE module shall define all task bodies created by the RTE generator.

Note that in vendor mode it is assumed that greater knowledge of the OS is available and therefore the above requirement does *not* apply so that specific optimizations, such as creating each task in a separate module, can be applied.

5.3.7.3.5 Lifecycle API

[rte_sws_1197] The RTE module shall define the RTE lifecycle API.

The RTE lifecycle API is defined in Section 5.8.

5.3.7.4 Reentrancy

All code invoked by generated RTE code that can be subject to concurrent execution must be reentrant. This requirement for reentrancy can be overridden if the generated code is not subject to concurrent execution, for example, if protected by a data consistency mechanism to ensure that access to critical regions is serialized.

5.4 RTE Data Structures

Object-code software components are compiled against an application header file created during the “RTE Contract” phase but are linked against an RTE (and application header file) created during the “RTE Generation” phase. When generated in compatibility mode, an RTE has to work for object-code components compiled against an application header file created in compatibility mode, even if the application header file was created by a different RTE generator. It is thus necessary to define the data structures and naming conventions for the compatibility mode to ensure that the object-code is compatible with the generated RTE. An RTE generated in vendor mode only has to work for those object-code components that were compiled against application header files created in vendor mode by a compatible RTE generator (which in general would mean an RTE generator supplied by the same vendor).

The use of standardized data structures imposes tight constraints on the RTE implementation and therefore restricts the freedom of RTE vendors to optimize the solution of object-code components but has the advantage that RTE generators from different vendors can be used to compile an object-code software-component and to generate the RTE. No such restrictions apply for the vendor mode. If an RTE generator operating in vendor mode is used for an object-code component in both phases, vendor-specific optimizations can be used.

Note that with the exception of data structures required for support object-code software components in compatibility mode, the data structures used for “RTE Generation” phase are not defined. This permits vendor specific API mappings and data structures to be used for a generated RTE without loss of portability.

The following definitions only apply to RTE generators operating in compatibility mode – in this mode the instance handle and the component data structure have to be defined even for those (object-code) software components for which multiple instantiation is forbidden to ensure compatibility.

5.4.1 Instance Handle

The RTE is required to support object-code components as well as multiple instances of the same AUTOSAR software-component mapped to an ECU [RTE00011]. To minimise memory overhead all instances of a component on an ECU share code

[RTE00012] and therefore both the RTE and the component instances require a means to distinguish different instances.

Support for both object-code components and multiple instances requires a level of indirection so that the correct generated RTE custom function is invoked in response to a component action. The indirection is supplied by the instance handle in combination with the API mapping defined in Section 5.2.6.

[rte_sws_1012] The component instance handle shall identify particular instances of a component.

The instance handle is passed to each runnable entity in a component when it is activated by the RTE as the first parameter of the function implementing the runnable entity `rte_sws_1016`. The instance handle is then passed back by the runnable entity to the RTE, as the first parameter of each direct RTE API call, so that the RTE can identify the correct component instance making the call. This scheme permits multiple instances of a component on the same ECU to share code.

The instance handle indirection permits the name of the RTE API call that is used within the component to be unique within the scope of a component as well as independent of the component's instance name. It thus enables object-code AUTOSAR software-components to be compiled before the final "RTE Generation" phase when the instance name is fixed.

[rte_sws_1013] For the RTE C/C++ API, any call that can operate on different instances of a component that supports multiple instantiation `rte_sws_in_0004` shall have an instance handle as the first formal parameter.

[rte_sws_3806] If a component does not support multiple instantiation, the instance handle parameter shall be omitted in the RTE C/C++ API and in the signature of the RTE Hook functions.

If the component does not support multiple instantiation, the name of the instance handle must be specified, since it is not passed to the API calls and runnable entities as parameters.

[rte_sws_3793] If a software component does not support multiple instantiation, the name of the instance handle shall be `Rte_Inst_<c>`, where `<c>` is the component type name.

The data type of the instance handle is defined in Section 5.5.2.

5.4.2 Component Data Structure

Different component instances share many common features - not least of which is support for shared code. However, each instance is required to invoke different RTE API functions and therefore the instance handle is used to access the component data structure that defines all instance specific data.

It is necessary to define the component data structure to ensure compatibility between the two RTE phases when operating in compatibility mode – for example, a “clever” compiler and linker may encode type information into a pointer type to ensure type-safety. In addition, the structure definition cannot be empty since this is an error in ANSI C.

[rte_sws_3713] The component data structure type shall be defined in the AUTOSAR Types Header file.

[rte_sws_3714] The type name of the component data structure shall be `Rte_CDS_<c>` where `<c>` is the component type name.

The members of the component data structure include function pointers. It is important that such members are not subject to run-time modification and therefore the component data structure is required to be placed in read-only memory.

[rte_sws_3715] All instances of the component data structure shall be defined as “const” (i.e. placed in read-only memory).

The elements of the component data structure are sorted into sections, each of which defines a logically related section. The sections defined within the component data structure are:

- **[rte_sws_3718]** Data Handles section.
- **[rte_sws_3719]** Per-instance Memory Handles section.
- **[rte_sws_1349]** Inter-runnable Variable Handles section.
- **[rte_sws_3720]** Calibration Parameter Handles section.
- **[rte_sws_3721]** Exclusive-area Handles section.
- **[rte_sws_3716]** Port API section.
- **[rte_sws_3717]** Inter Runnable Variable API section.
- **[rte_sws_3722]** Vendor specific section.

The order of elements within each section of the component data structure is defined as follows;

[rte_sws_3723] Section entries shall be sorted alphabetically (ASCII / ISO 8859-1 code in ascending order) unless stated otherwise.

The sorting of entries is applied to each section in turn.

Note that there is *no* prefix associated with the name of each entry within a section; the component data structure as a whole has the prefix and therefore there is no need for each member to have the same prefix.

ANSI C does not permit empty structure definitions yet an instance handle is required for the RTE to function. Therefore if there are no API calls then a single dummy entry is defined for the RTE.

[rte_sws_3724] If all sections of the Component Data Structure are empty the Component Data Structure shall contain a `uint8` with name `_dummy`.

5.4.2.1 Data Handles Section

The data handles section is required to support the `Rte_IRead` and `Rte_IWrite` calls (see Section 5.2.4).

[rte_sws_3733] Data Handles shall be named `<re>_<p>_<d>` where `<re>` is the runnable entity name that reads (or writes) the data item, `<p>` the port name, `<d>` the data element.

A runnable cannot read *and* write to the same port/data element since the port is inherently uni-directional (a provide port can only be written, a required port can only be read).

[rte_sws_2608] The Data Handle shall be a pointer to a `Data Element` with `Status` if and only if the runnable has read access and either

- `data element outdated notification` or
- `data element invalidation`

is activated for this `data element`.

[rte_sws_2588] Otherwise, the data type for a Data Handle shall be a pointer to either a `Data Element` without `Status`.

See below for the definitions of these terms.

5.4.2.1.1 Data Element without Status

[rte_sws_1363] The data type for a “Data Element without Status” shall be named `Rte_DE_<dt>` where `<dt>` is the data element type.

[rte_sws_1364] A `Data Element without Status` shall be a structure containing a single member named `value`.

[rte_sws_2607] The `value` member of a `Data Element without Status` shall have the same data type as the corresponding `DataElement`.

Note that requirements `rte_sws_1364` and `rte_sws_2607` together imply that creating a variable of data type `Rte_DE_<dt>` allocates enough memory to store the data copy.

5.4.2.1.2 Data Element with Status

[rte_sws_1365] The data type for a “Data Element with Status” shall be named `Rte_DES_<dt>` where `<dt>` is the data element type.

[rte_sws_1366] A Data Element with Status shall be a structure containing two members.

[rte_sws_3734] The first member of each Data Element with Status shall be named 'value'

[rte_sws_2666] The value member of a Data Element with Status shall have the type of the corresponding DataElement.

[rte_sws_2589] The second member of each Data Element with Status shall be named 'status'.

[rte_sws_2590] The status member of a Data Element with Status shall be of the Std_ReturnType type.

[rte_sws_2609] The status member of a Data Element with Status shall contain the error status corresponding to the value member.

5.4.2.1.3 Usage

[rte_sws_1367] A definition for every required Data Element with Status and every Data Element without Status must be emitted in the *AUTOSAR Types Header File*.

The AUTOSAR Types Header File is defined in Section 5.3.4).

Example 5.16

Consider a uint8 data element, a, of port p which is accessed using DataWriteAccess semantics by runnables re1 and re2 and DataReadAccess semantics by runnable re2 within component c. data element outdated is defined for this DataElementPrototype.

The required data types within the *AUTOSAR Types Header File* would be:

```

1 typedef struct {
2     uint8 value;
3 } Rte_DE_uint8;
4
5 typedef struct {
6     uint8 value;
7     Std_ReturnType status;
8 } Rte_DES_uint8;
```

The component data structure for c would also include:

```

1 Rte_DE_uint8* re1_p_a;
2 Rte_DES_uint8* re2_p_a;
```

A software-component that is supplied as object-code or is multiply instantiated requires “general purpose” definitions of Rte_IRead, Rte_IWrite, and

`Rte_IStatus` that use the data handles to access the data copies created within the generated RTE. For example:

```

1 #define Rte_IWrite_re1_p_a(s,v) ((s)->re1_p_a->value = (v))
2 #define Rte_IWrite_re2_p_a(s,v) ((s)->re2_p_a->value = (v))
3 #define Rte_IRead_re2_p_a(s,v) ((s)->re2_p_a->value)
4 #define Rte_IStatus_re2_p_a(s) ((s)->re2_p_a->status)

```

The definitions of `Rte_IRead`, `Rte_IWrite`, and `Rte_IStatus` are type-safe since an attempt to assign an incorrect type will be detected by the compiler.

For source code component that does **not** use multiple instantiation the definitions of `Rte_IRead`, `Rte_IWrite`, and `Rte_IStatus` can remain as above or vendor specific optimizations can be applied without loss of portability.

The values assigned to data handles within *instances* of the component data structure created within the generated RTE depend on the mapping of tasks and runnables – See Section 5.2.4.

5.4.2.2 Per-instance Memory Handles Section

The Per-instance Memory Section Handles section enables to access instance specific memory (sections).

[rte_sws_2301] The CDS shall contain a handle for each Per-instance Memory. This handle member shall be named `Pim_<name>` where `<name>` is the per-instance memory name.

The Per-instance Memory Handles are typed;

[rte_sws_2302] The data type of each Per-instance Memory Handle shall be a pointer to the type of the per instance memory that is defined in the RTE Types header file.

The RTE supports the access to the per-instance memories by the `Rte_Pim` API.

Example 5.17

Referring to the specification items `rte_sws_2301` and `rte_sws_2302` Example 5.4 can be extended –

with respect to the software-component header:

```

1 struct Rte_CDS_c {
2     ...
3     /* per-instance memory handle section */
4     Rte_PimType_c_MyMemType *Pim_mem;
5
6     ...
7 };
8
9 #define Rte_Pim_mem(s) ((s)->Pim_mem)

```

and in Rte.c:

```

1  Rte_PimType_c_MyMemType mem1;
2
3  const struct Rte_CDS_c Rte_Instance_c1 = {
4      ...
5      /* per-instance memory handle section */
6      /* Rte_PimType_c_MyMemType Pim_mem */
7      &mem1
8      ...
9  };

```

5.4.2.3 Inter Runnable Variable Handles Section

Each runnable may require separate handling for the inter runnable variables that it accesses. The indirection required for explicit access to inter runnable variables is described in section 5.4.2.7. The inter runnable variable handles section within the component data structure contains pointers to the (shadow) memory of inter runnable variables that can be directly accessed with the implicit API macros. The inter runnable variable handles section does not contain pointers for memory to handle inter runnable variables that are accessed with explicit API only.

[rte_sws_2636] For each runnable and each inter runnable variable that is accessed implicitly by the runnable, there shall be exactly one inter runnable handle member within the component data structure and this inter runnable variable handle shall point to the (shadow) memory of the inter runnable variable for the runnable.

[rte_sws_1350] The name of each inter runnable variable handle member within the component data structure shall be `Irv_<re>_<name>` where `<name>` is the Inter-Runnable Variable short name and `<re>` is short name of the runnable name.

[rte_sws_1351] The data type of each inter runnable variable handle member shall be a pointer to the type of the inter runnable variable.

5.4.2.4 Exclusive-area handles Section

The exclusive area handles section includes exclusive areas that are accessed explicitly, using the RTE API, by the SW-C. Each entry in the section is a function pointer to the relevant RTE API function generated for the SW-C instance.

[rte_sws_3739] The name of each Exclusive-area handle shall be `<root>_<name>` where `<root>` is either `Entry` or `Exit` and `<name>` is the Exclusive-area name.

[rte_sws_3740] The data type of each Exclusive-area handle entry shall be a function pointer that points to the generated RTE API function.

[rte_sws_3812] Entries in the Exclusive-area handles section shall be sorted alphabetically.

Note that two function pointers will be required for each accessed exclusive area; one for the Entry function and one for the Exit function.

5.4.2.5 Port API Section

Port API section comprises zero or more *function references* within the component data structure type that defines all API functions that access a port and can be invoked by the software-component (instance).

[rte_sws_2616] The function table entries for port access shall be grouped by the port names into port data structures.

Each entry in the port API section of the component data structure is a “port data structure”.

[rte_sws_2617] The name of each port data structure in the component data structure shall be `<p>` where `<p>` is the port short-name.

[rte_sws_3799] The component data structure shall contain a port data structure for port `p` only if the component supports multiple instantiation or if the **indirectAPI** attribute for `p` is set to 'true'.

[rte_sws_3731] The data type name for a port data structure shall be `struct Rte_PDS_<c>_<i>_<P/R>` where `<c>` is the component type name, `<i>` is the port interface name and 'P' or 'R' are literals to indicate provide or require ports respectively.

[rte_sws_3732] The port data structure type(s) shall be defined in the AUTOSAR types header file.

A port data structure type is defined for each port interface that types a port. Thus different ports typed by the same port interface structure share the same port data structure type.

[rte_sws_3601] The AUTOSAR types header file shall contain a definition of a port data structure type for interface `i` and port type `R` or `P` only if the component supports multiple instantiation or at least one require or provide port exists that has the **indirectAPI** attribute set to 'true'.

[rte_sws_3730] A port data structure shall contain a function table entry for each API function associated with the port as referenced in table 5.3. Pure API macros, like `Rte_IRead` and other implicit API functions, do not have a function table entry.

API function	reference
<code>Rte_Send_<p>_<d></code>	5.6.4
<code>Rte_Write_<p>_<d></code>	5.6.4
<code>Rte_Switch_<p>_<m></code>	5.6.5
<code>Rte_Invalidate_<p>_<d></code>	5.6.6
<code>Rte_Feedback_<p>_<d></code>	5.6.7

API function	reference
Rte_Read_<p>_<d>	5.6.8
Rte_Receive_<p>_<d>	5.6.9
Rte_Call_<p>_<o>	5.6.10
Rte_Result_<p>_<o>	5.6.11
Rte_Calprm_<p>_<name>	5.6.14
Rte_Mode_<p>_<o>	5.6.26

Table 5.3: Table of API functions that are referenced in the port API section.

[rte_sws_2620] An API function shall only be included in a port data structure, if it is required at least by one port.

[rte_sws_2621] If a function table entry is available in a port data structure, the corresponding function shall be implemented for all ports that use this port data structure type. API functions related to ports that are not required by the AUTOSAR configuration shall behave like those for an unconnected port.

APIs may be required only for some ports of a software component instance due to differences in for example the need for transmission acknowledgement. `rte_sws_2621` is necessary for the concept of the indirect API. It allows iteration over ports.

[rte_sws_ext_3813] The indirect API may only be used for a port if explicit data access to this port is specified via `DataSendPoints` and `DataReceivePoints`.

[rte_sws_1055] The name of each function table entry in a port data structure shall be `<name>_<d/o>` where `<name>` is the API root (e.g. `Call`, `Write`) and `<d/o>` the data element or operation name.

Requirement `rte_sws_1055` does *not* include the port name in the function table entry name since the port is implicit when using a port handle.

[rte_sws_3726] The data type of each function table entry in a port data structure shall be a function pointer that points to the generated RTE function.

The signature of a generated function, and hence the definition of the function pointer type, is the same as the signature of the relevant RTE API call (see Section 5.6) with the exception that the instance handle is omitted.

Example 5.18

This example shows a port data structure for the provide ports of the interface type `i2` in an AUTOSAR SW-C `c`.

`i2` is a `SenderReceiverInterface` which contains a data element prototype of type `uint8` with `isQueued` set to `false`.

If one of the provide ports of `c` for the interface `i2` has a transmission acknowledgement defined and `i2` is not used with `data element`

invalidation, the AUTOSAR types header file would include a port data structure type like this:

```

1 struct Rte_PDS_c_i2_P {
2     Std_ReturnType (*Feedback_a) (uint8);
3     Std_ReturnType (*Write_a) (uint8);
4 }

```

If the provide port `p1` of the AUTOSAR SW-C `c` is of interface `i2`, the generated component header file would include the following macros to provide the direct API functions `Rte_Feedback_p1_a` and `Rte_Write_p1_a`:

```

1 /*direct API*/
2 #define Rte_Feedback_p1_a(inst, data)
3     ((inst)->p1.Feedback_a) (data)
4 #define Rte_Write_p1_a(inst, data) ((inst)->p1.Write_a) (data)

```

[rte_sws_2618] The port data structures within a component data structure shall first be sorted on the port data structure type name and then on the short name of the port.

The requirements `rte_sws_3731` and `rte_sws_2618` guarantee, that all port data structures within the component data structure are grouped by their interface type and require/provide-direction.

Example 5.19

This example shows the grouping of port data structures within the component data structure.

The AUTOSAR types header file for an AUTOSAR SW-C `c` with three provide ports `p1`, `p2`, and `p3` of interface `i2` would include a block of port data structures like this in the generated AUTOSAR Types Header file:

```

1 struct Rte_CDS_c {
2     ...
3     struct Rte_PDS_c_i1_R z;
4
5     /* component data structures          *
6     * for provide ports of interface i2 */
7     struct Rte_PDS_c_i2_P p1;
8     struct Rte_PDS_c_i2_P p2;
9     struct Rte_PDS_c_i2_P p3;
10
11    /*further component data structures*/
12    struct Rte_PDS_c_i2_R c;
13    ...
14 }
15

```

If `inst` is a pointer to a component data structure, and `ph` is defined by

```

1 struct Rte_PDS_c_i2_P *ph = &(inst->p1);

```

`ph` points to the port data structure `p1` of the instance handle `inst`. Since the three provide port data structures `p1`, `p2`, and `p3` of interface `i2` are ordered sequentially in the component data structure, `ph` can also be interpreted as an array of port data structures. E.g., `ph[2]` is equal to `inst->p3`.

In the following, `ph` will be called a port handle.

[rte_sws_1343] RTE shall create port handle types for each port data structure using `typedef` to a pointer to the appropriate port data structure.

[rte_sws_1342] The port handle type name shall be `Rte_PortHandle_<i>_<P/R>` where `<i>` is the port interface name and 'P' or 'R' are literals to indicate provide or receive ports respectively.

[rte_sws_1053] The port handle types shall be written to the application header file.

The port handle types cannot be included in the AUTOSAR types header file due to potential name clashes between components.

RTE provides port handles for access to the arrays of port data structures of the same interface type and provide/receive direction by the macro `Rte_Ports`, see section 5.6.1, and to the number of similar ports by the macro `Rte_NPorts`, see 5.6.1.

Example 5.20

For the provide port `i2` of AUTOSAR SW-C `c` from example 5.18, the following port handle type will be defined in the component header file:

```
1 typedef struct Rte_PDS_c_i2_P *Rte_PortHandle_i2_P;
```

The macros to access the port handles for the indirect API might look like this in the generated component header file:

```
1 /*indirect (port oriented) API*/
2 #define Rte_Ports_i2_P(inst) &((inst)->p1)
3 #define Rte_NPorts_i2_P(inst) 3
```

So, the port handle `ph` of the previous example 5.19 could be defined by a user as:

```
1 Rte_PortHandle_i2_P ph = Rte_Ports_i2_P(inst);
```

To write '49' on all ports `p1` to `p3`, the indirect API can be used within the software component as follows:

```
1 uint8 p;
2 Rte_PortHandle_i2_P ph = Rte_Ports_i2_P(inst);
3 for (p=0;p<Rte_NPorts_i_P(inst);p++) {
4     ph[p].Write_a(49);
5 }
```

Software components may also want to set up their own port handle arrays to iterate over a smaller sub group than all ports with the same interface and direction. `Rte_Port`

can be used to pick the port handle for one specific port, see 5.6.3.

5.4.2.6 Calibration Parameter Handles Section

The RTE is required to support access to calibration parameters derived by *per-instance* CalprmElementPrototypes (see 4.2.6.3) using the `Rte_CData` (see section 5.6.13).

[rte_sws_6029] The name of each Calibration parameter handle shall be `CData_<name>` where `<name>` is the CalprmElementPrototype name.

[rte_sws_3949] The type of Each calibration parameter handle shall be a function pointer that points to the generated RTE function.

The function pointer points to the generated RTE function and therefore the return value of the function call depends on the data type of the CalprmComponentPrototype; it is the value for primitive data types whereas a reference is returned for complex data types.

Note that accesses to CalprmElementPrototypes within CalprmComponentTypes do not require handles within this section since the generated `Rte_Calprm` (see section 5.6.14) API is accessed either directly (single instantiation) or through handles in the port API section (multiple instantiation). Likewise, access to *shared* CalprmElement-Prototypes does not require a handle since, by definition, no per-instance data is present.

5.4.2.7 Inter Runnable Variable API Section

The Inter Runnable Variable API section comprises zero or more *function table entries* within the component data structure type that defines all explicit API functions to access an inter runnable variable by the software-component (instance). The API for implicit access of inter runnable variables does not have any *function table entries*, since the implicit API uses macro's to access the inter runnable variables or their shadow memory directly, see section 5.4.2.3.

Since the entries of this section are only required to access the explicit InterRunnable-Variable API if a software component supports multiple instantiation, it shall be omitted for software components which do not support multiple instantiation.

[rte_sws_3725] If the component supports multiple instantiation, the member name of each function table entry within the component data structure shall be `<name>_<re>_<d>` where `<name>` is the API root (e.g. `IrvRead`), `<re>` the runnable name, and `<d>` the inter runnable variable name.

[rte_sws_3752] The data type of each function table entry shall be a function pointer that points to the generated RTE function.

The signature of a generated function, and hence the definition of the function pointer type, is the same as the signature of the relevant RTE API call (see Section 5.6) with the exception that the instance handle is omitted.

[rte_sws_2623] If the component supports multiple instantiation, the inter runnable variable API section shall contain pointers to the following API functions:

API function	reference
Rte_IrvRead_<re>_<d>	5.6.22
Rte_IrvWrite_<re>_<d>	5.6.23

Table 5.4: Table of API functions that are referenced in the inter runnable variable API section

[rte_sws_3791] If the software component does not support multiple instantiation, the inter runnable variable API section shall be empty.

5.4.2.8 Vendor Specific Section

The vendor specific section is used to contain any vendor specific data required to be supported for each instances. By definition the contents of this section are outside the scope of this chapter and only available for use by the RTE generator responsible for the “RTE Generation” phase.

5.5 API Data Types

Besides the API functions for accessing RTE services, the API also contains RTE-specific data types.

5.5.1 Std_ReturnType

The specification in [12] specifies a standard API return type `Std_ReturnType`. The `Std_ReturnType` defines the “status” and “error values” returned by API functions. It is defined as a `uint8` type. The value “0” is reserved for “No error occurred”.

Figure 5.4 shows the general layout of `Std_ReturnType`.

The two most significant bits of the `Std_ReturnType` are reserved flags:

- The most significant bit 7 of `Std_ReturnType` is the “Immediate Infrastructure Error Flag” with the following values
 - “1” the error code indicates an immediate infrastructure error.
 - “0” the error code indicates no immediate infrastructure error.

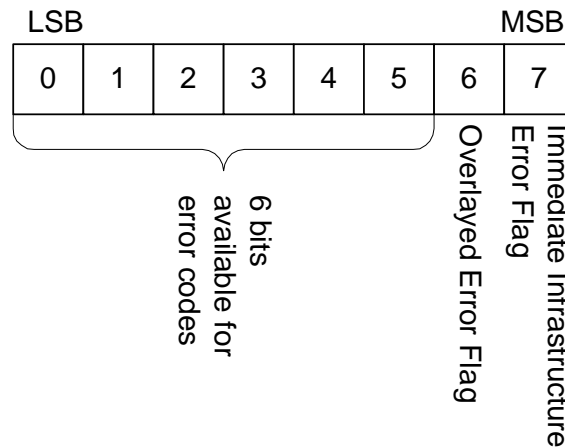


Figure 5.4: Bit-Layout of the Std_ReturnType

- The second most significant bit 6 of `Std_ReturnType` is the Overlaid Error Flag. The use of this flag depends on the context and will be explained in table 5.6.

5.5.1.1 Infrastructure Errors

Infrastructure errors are split into two groups:

- “Immediate Infrastructure Errors” can be associated with the currently available data set. These Immediate Infrastructure Errors are mutually exclusive. Only one of these errors can be notified to a SW-C with one API call.

[rte_sws_2593] Immediate Infrastructure Errors shall override any application level error.

Immediate Infrastructure Error codes are used on the receiver side for errors that result in no reception of application data and application errors.

An Immediate Infrastructure Error is indicated in the Std_ReturnType by the Immediate Infrastructure Error Flag being set.

- “Overlaid Errors” are associated with communication events that happened after the reception of the currently available data set, e.g., data element outdated notification, or loss of data elements due to queue overflow.

[rte_sws_1318] Overlaid Error Flags shall be reported using the unique bit of the Overlaid Error Flag within the Std_ReturnType type.

An Overlaid Error can be combined with any other application or infrastructure error code.

5.5.1.2 Application Errors

[rte_sws_2573] RTE shall support application errors with the following format definition: Application errors are coded in the least significant 6 bits of Std_ReturnType with the Immediate Infrastructure Error Flag set to “0”. The application error code does not use the Overlaid Error Flag.

This results in the following value range for application errors:

range	minimum value	maximum value
application errors	1	63

Table 5.5: application error value range

In client server communication, the server may return any value within the application error range. The client will then receive one of the following:

- An Immediate Infrastructure Error to indicate that the communication was not successful or
- The server return code or

- The server return code might be overlaid by the Overlaid Error Flag in a future release of RTE. In this release, there is no overlaid error defined for client server communication.

The client can filter the return value, e.g., by using the following code:

```
Std_ReturnType status;
status = Rte_Call_<p>_<d>(<instance>, <parameters>);
if (status & 64) {
    /* handle overlaid error flag *
    * in this release of the RTE, the flag is reserved *
    * but not used for client server communication */
}
status &= (Std_ReturnType) (~64);
if(status & 128) {
    /* handle infrastructure error */
}
else {
    /* handle application error with error code status */
}
```

5.5.1.3 Predefined Error Codes

[rte_sws.in_2622] For client server communication, application error values are defined per client server interface and shall be passed to the RTE with the interface configuration.

The following standard error and status identifiers are defined:

Symbolic name	Value	Comments
[rte_sws_1058] RTE_E_OK	0	No error occurred.

Standard Application Error Values:		
[rte_sws_2594] RTE_E_INVALID	1	Generic application error indicated by signal invalidation in sender receiver communication with isQueued = false on the receiver side.
To be defined by the corresponding AUTOSAR Service	1	Returned by AUTOSAR Services to indicate a generic application error.

Immediate Infrastructure Error codes

Symbolic name	Value	Comments
[rte_sws_1060] RTE_E_COM_STOPPED	128	An IPDU group was disabled while the application was waiting for the transmission acknowledgment. No value is available. This is not considered a fault, since the IPDU group is switched off on purpose. This semantics are as follows: <ul style="list-style-type: none"> • The OUT buffers of a client or of explicit read APIs are not modified • no runnable with startOnEvent on a DataReceivedEvent for this dataElement-Prototype is triggered. • the buffers for implicit read access will keep the previous value.
[rte_sws_1064] RTE_E_TIMEOUT	129	A blocking API call returned due to expiry of a local timeout rather than the intended result. OUT buffers are not modified. The interpretation of this being an error depends on the application.
[rte_sws_1317] RTE_E_LIMIT	130	A internal RTE limit has been exceeded. Request could not be handled. OUT buffers are not modified.
[rte_sws_1061] RTE_E_NO_DATA	131	An explicit read API call returned no data. (This is no error.)
[rte_sws_1065] RTE_E_TRANSMIT_ACK	132	Transmission acknowledgement received.

Overlaid Errors		
These errors do not refer to the data returned with the API. They can be overlaid with other Application- or Immediate Infrastructure Errors.		
[rte_sws_2571] RTE_E_LOST_DATA	64	An API call for reading received data of isQueued = true indicates that some incoming data has been lost due to an overflow of the receive queue or due to an error of the underlying communication stack.
[rte_sws_2702] RTE_E_MAX_AGE_EXCEEDED	64	An API call for reading received data of isQueued = false indicates that the available data has exceeded the aliveTimeout limit. A COM signal outdated callback will result in this error.

Table 5.6: RTE Error and Status values

The underlying type for `Std_ReturnType` is defined as a `uint8` for reasons of compatibility – it avoids RTEs from different vendors assuming a different size if an `enum` was the underlying type. Consequently, `#define` is used to declare the error values:

```
1 typedef uint8 Std_ReturnType;
2
3 #define RTE_E_OK ((Std_ReturnType) 0)
```

[rte_sws_1269] The standard errors as defined in table 5.6 including `RTE_E_OK` shall be defined in the RTE Header File.

[rte_sws_2575] Application Error Identifiers with exception of `RTE_E_INVALID` shall be defined in the Application Header File.

[rte_sws_2576] The application errors shall have a symbolic name defined as follows:

```
1 #define RTE_E_<interface>_<error> <error value>
```

where `<interface>` `rte_sws_in_1352` and `<error>` `rte_sws_in_2574` are the interface and error names from the configuration.

An `Std_ReturnType` value can be directly compared (for equality) with the above pre-defined error identifiers.

5.5.2 Rte_Instance

The `Rte_Instance` data type defines the handle used to access instance specific information from the component data structure.

[rte_sws_1148] The underlying data type for an instance handle shall be a pointer to a *Component Data Structure*.

The component data structure (see Section 5.4.2) is uniquely defined for a component type and therefore the data type for the instance handle is automatically unique for each component type.

The instance handle type is defined in the application header file `rte_sws_1007`.

To avoid long and complex type names within SW-C code the following requirement imposes a fixed name on the instance handle data type.

[rte_sws_1150] The name of the instance handle type shall be defined, using `typedef` as `Rte_Instance`.

5.5.3 RTE Modes

An `Rte_ModeType` is used to hold the identifiers for the `ModeDeclarations` of a `ModeDeclarationGroup`.

[rte_sws_2627] For each ModeDeclarationGroup, the AUTOSAR Types HeaderFile shall contain a type definition

```
1 typedef <type> Rte_ModeType_<ModeDeclarationGroup>;
```

where <ModeDeclarationGroup> is the short name of the ModeDeclarationGroup and <type> is `uint8` for ModeDeclarationGroups with less than 256 ModeDeclarations and `uint16` for ModeDeclarationGroups with 256 or more ModeDeclarations.

Within the `Rte_ModeType_<ModeDeclarationGroup>`, the null value is reserved to express a transition between modes.

[rte_sws_2659] For each ModeDeclarationGroup, the AUTOSAR Types HeaderFile shall contain a definition

```
1 #define RTE_TRANSITION_<ModeDeclarationGroup> \  
2     ((Rte_ModeType_<ModeDeclarationGroup>)<n>)
```

where <ModeDeclarationGroup> is the short name of the ModeDeclarationGroup and <n> is the number of modes declared within the group.

[rte_sws_2568] For each mode of a mode declaration, the AUTOSAR Types Header file shall contain a definition

```
1 #define RTE_MODE_<ModeDeclarationGroup>_<ModeDeclaration> \  
2     ((Rte_ModeType_<ModeDeclarationGroup>)<index>)
```

where <ModeDeclarationGroup> is the short name of the ModeDeclarationGroup, <ModeDeclaration> is the short name of a ModeDeclaration, and <index> is the index of the ModeDeclarations in alphabetic ordering (ASCII / ISO 8859-1 code in ascending order) of the short names within the ModeDeclarationGroup. The lowest index shall be '0' and therefore the range of assigned values is 0 . . <n> where <n> is the number of modes declared within the group

5.5.4 Enumeration Data Types

Enumeration is not a plain primitive data type. Rather a range of integers can be used as a structural description. The mapping of integers on "labels" in the enumeration is actually modelled in the SwC-T with the semantics class `CompuMethod` of a `SwDataDefProps` [17]. Enumeration data types are modeled as `PrimitiveTypes` having a `SwDataDefProps` referencing a `CompuMethod` that contains only `CompuScales` with point ranges (i. e. lower and upper limit of a `CompuScale` are identical).

[rte_sws_3809] The AUTOSAR Types header file shall include the definitions of all enumeration constants of AUTOSAR data types irrespective of their use by the generated RTE.

This requirement ensures the availability of AUTOSAR data type enumeration constants for the internal use in AUTOSAR software components.

[rte_sws_3810] For each CompuScale of a CompuMethod with category "TEXTTABLE" that contains only CompuScales with point ranges the AUTOSAR Types header file shall contain a definition

```
1 #ifndef <EnumLiteral>
2 #define <EnumLiteral> ((<type>) <value>)
3 #endif /* <EnumLiteral> */
```

where `<EnumLiteral>` is the string specified in the VT element of the CompuConst of the respective CompuScale. `<type>` is the identifier of the PrimitiveType the CompuMethod belongs to. `<value>` is the value representing the CompuScale's point range.

rte_sws_3810 implies that the RTE does not add any prefix to the names of the enumeration constants. This is necessary in order to handle enumeration constants supplied by Basic Software modules which all use their own prefix convention. Enumeration constant names have to be unique in the whole AUTOSAR system.

[rte_sws_3813] If the input of the RTE generator contains two or more CompuMethods with category "TEXTTABLE" that both contain a CompuScale with a point range and an identical enumeration literal name specified in the VT element of the CompuScale, the RTE generator shall reject this input as an invalid configuration.

5.5.5 Range Data Types

For the AUTOSAR data types `IntegerType` and `RealType` a `Range` has to be specified giving the `lowerLimit` and the `upperLimit`. To allow a Software Component the access to these values two definitions for these values shall be generated.

[rte_sws_5051] The AUTOSAR Types header file shall include the definitions of all `Range` constants of AUTOSAR data types irrespective of their use by the generated RTE.

[rte_sws_5052] For AUTOSAR data types which inherit from `Range` the AUTOSAR Types header file shall contain two definitions

```
1 #define <DataType>_LowerLimit ((<DataType>) <lowerLimitValue>)
2 #define <DataType>_UpperLimit ((<DataType>) <upperLimitValue>)
```

where `<DataType>` is the short name of the data type. `<lowerLimitValue>` and `<upperLimitValue>` are the values of the respective range.

5.6 API Reference

The functions described in this section are organized by the RTE API mapping name used by C and C++ AUTOSAR software-components to access the API. The API mapping hides from the AUTOSAR software-component programmer any need to be aware

of the steps taken by the RTE generator to ensure that the generated API functions have unique names.

The instance handle as the first parameter of the API calls is marked as an optional parameter in this section. If an AUTOSAR software-component supports multiple instantiation, the instance handle shall be passed `rte_sws_1013`.

Note that `rte_sws_3806` requires that the instance handle parameter does not exist if the AUTOSAR software-component does not support multiple instantiation.

5.6.1 Rte.Ports

Purpose: Provide an array of the ports of a given interface type and a given provide / require usage that can be accessed by the indirect API.

Signature: **[rte_sws_2619]**
`Rte_PortHandle_<i>_<R/P>`
`Rte_Ports_<i>_<R/P> ([IN Rte_Instance])`

Where here `<i>` is the port interface name and 'P' or 'R' are literals to indicate provide or require ports respectively.

Existence: **[rte_sws_2613]** An `Rte_Ports` API shall be created for each interface type and usage by a port when the **indirectAPI** attribute of the port is set to true.

Description: The `Rte_Ports` API provides access to an array of ports for the port oriented API.

[rte_sws_3602] Only those ports for which the indirect API was generated shall be contained in the array of ports.

Return Value: Array of port data structures of the corresponding interface type and usage.

Notes: None.

5.6.2 Rte.NPorts

Purpose: Provide the number of ports of a given interface type and provide / require usage that can be accessed through the indirect API.

Signature: **[rte_sws_2614]**
`uint8`
`Rte_NPorts_<i>_<R/P> ([IN Rte_Instance])`

Where here `<i>` is the port interface name and 'P' or 'R' are literals to indicate provide or require ports respectively.

- Existence:** [rte_sws_2615] An `Rte_NPorts` API shall be created for each interface type and usage by a port when the **indirectAPI** attribute of the port is set to true.
- Description:** The `Rte_NPorts` API supports access to an array of ports for the port oriented API.
- [rte_sws_3603] The `Rte_NPorts` shall return only the number of ports of a given interface and provide / require usage for which the indirect API was generated.
- Return Value:** Number of port data structures of the corresponding interface type and usage.
- Notes:** None.

5.6.3 Rte_Port

- Purpose:** Provide access to the port data structure for a single port of a particular software component instance. This allows a software component to extract a sub-group of ports characterized by the same interface in order to iterate over this sub-group.
- Signature:** [rte_sws_1354]
`Rte_PortHandle_<i>_<R/P>`
`Rte_Port_<p>([IN Rte_Instance])`
 where <i> is the port interface name and <p> is the name of the port.
- Existence:** [rte_sws_1355] The `Rte_Port` API shall be created for each port of an AUTOSAR SW-C for which the **indirectAPI** attribute is set to true.
- Description:** The `Rte_Port` API provides a pointer to a single port data structure, in order to support the indirect API.
- Return Value:** Pointer to port data structure for the appropriate port.
- Notes:** None.

5.6.4 Rte_Send/Rte_Write

- Purpose:** Initiate an “explicit” sender-receiver transmission of data elements. The `Rte_Write` API call is used for “data” (`isQueued = false`) and the `Rte_Send` API call used for “events” (`isQueued = true`).
- Signature:** [rte_sws_1071]

```
Std_ReturnType
Rte_Write_<p>_<o>([IN Rte_Instance <instance>],
                 IN <data>)
```

[rte_sws_1072]

```
Std_ReturnType
Rte_Send_<p>_<o>([IN Rte_Instance <instance>],
                IN <data>)
```

Where <p> is the port name and <o> the DataElementPrototype within the sender-receiver interface categorizing the port.

Existence: **[rte_sws_1280]** The presence of a DataSendPoint for a provided DataElementPrototype with isQueued = false shall result in the generation of an Rte_Write API for the provided DataElementPrototype.

[rte_sws_1281] The presence of a DataSendPoint for a provided DataElementPrototype with isQueued = true shall result in the generation of an Rte_Send API for the provided DataElementPrototype.

[rte_sws_ext_2680] The Rte_Send/Rte_Write APIs may only be used by the runnable that contains the corresponding DataSendPoint

Description: The Rte_Send and Rte_Write API calls initiate a sender-receiver communication where the transmission occurs at the point the API call is made (cf. explicit transmission).

The Rte_Send and Rte_Write API calls include exactly one IN parameter for the data element – this will be passed by value for primitive data types and by reference for all other types.

If the IN parameter is passed by reference, the pointer must remain valid until the API call returns.

Return Value: The return value is used to indicate errors detected by the RTE during execution of the APICallWrite or APICallSend.

- **[rte_sws_1073]** RTE_E_OK – data passed to communication service successfully.
- **[rte_sws_1074]** RTE_E_COM_STOPPED – the RTE could not perform the operation because the COM service is currently not available (inter ECU communication only). RTE shall return RTE_E_COM_STOPPED when the corresponding COM service returns COM_E_STOPPED.
- **[rte_sws_2634]** RTE_E_LIMIT – an ‘event’ has been discarded due to a full queue. (intra ECU communication only).

Notes: The Rte_Write and Rte_Send calls are closely related – Rte_Write is used to transmit “data” (isQueued = false) and Rte_Send to transmit “events” (isQueued = true).

[rte_sws_1077] In case of inter ECU communication, the `Rte_Write` and `Rte_Send` shall cause an immediate transmission request.

Note that depending on the configuration a transmission request may not result in an actual transmission, for example transmission may be rate limited (time-based filtering) and thus dependent on other factors than API calls.

[rte_sws_1081] In case of inter ECU communication, the `Rte_Write` or `Rte_Send` API shall return when the signal has been passed to the communication service for transmission.

Depending on the communication server the transmission may or may not have been acknowledged by the receiver at the point the API call returns.

[rte_sws_2633] In case of intra ECU communication, the `Rte_Send` API call shall return after attempting to enqueue the data.

[rte_sws_2635] In case of intra ECU communication, the `Rte_Write` API call shall return after copying the data.

[rte_sws_1080] If the transmission acknowledgement is enabled, the RTE shall notify component when the transmission is acknowledged or a transmission error occurs.

[rte_sws_1082] If a provide port typed by a sender-receiver interface has multiple require ports connected (i.e. it has multiple receivers), then the RTE shall ensure that writes to all receivers are independent.

Requirement `rte_sws_1082` ensures that an error detected by the RTE when writing to one receiver, e.g. an overflow in one component's queue, does not prevent the transmission of this message to other components.

5.6.5 Rte_Switch

Purpose: Initiate a mode switch. The `Rte_Switch` API call is used for 'explicit' sending of a mode switch notification.

Signature: **[rte_sws_2631]**
`Std_ReturnType`
`Rte_Switch_<p>_<o>([IN Rte_Instance <instance>],`
`IN Rte_ModeType_<M> <mode>)`

Where `<p>` is the port name and `<o>` the `ModeDeclarationGroup-Prototype` within the sender-receiver interface categorizing the port.

Existence: **[rte_sws_2632]** The existence of a `ModeSwitchPoint` shall result in the generation of a `Rte_Switch` API.

[rte_sws_ext_2681] The `Rte_Switch` API may only be used by the runnable that contains the corresponding `ModeSwitchPoint`

Description: The `Rte_Switch` triggers a synchronous mode switch for all connected require `ModeDeclarationGroupPrototypes`.

The `Rte_Switch` API call includes exactly one IN parameter for the next mode `<mode>` of type `Rte_ModeType_<M>` where `<M>` is the `ModeDeclarationGroup` short name.

Return Value: The return value is used to indicate errors detected by the RTE during execution of the `Rte_Switch` call.

- **[rte_sws_2674]** `RTE_E_OK` – data passed to service successfully.
- **[rte_sws_2675]** `RTE_E_LIMIT` – a mode switch has been discarded due to a full queue.

Notes: `Rte_Switch` is restricted to ECU local communication.

If a mode instance is currently involved in a transition then the `Rte_Switch` API will attempt to queue the request and return `rte_sws_2667`. However if no transition is in progress for the mode instance, the mode disablings and the activations of `OnEntry` and `OnExit` runnables for this mode instance are executed before the `Rte_Switch` API returns `rte_sws_2665`.

Note that the mode switch might be discarded when the queue is full and a mode transition is in progress, see `rte_sws_2675`.

[rte_sws_2673] If the mode switched acknowledgment is enabled, the RTE shall notify the mode manager when the mode switch is completed.

5.6.6 Rte_Invalidate

Purpose: Invalidate a data element for an “explicit” sender-receiver transmission.

Signature: **[rte_sws_1206]**
`Std_ReturnType`
`Rte_Invalidate_<p>_<o>([IN Rte_Instance <instance>])`

Where `<p>` is the port name and `<o>` the data element within the sender-receiver interface categorizing the port.

Existence: **[rte_sws_1282]** An `Rte_Invalidate` API shall be created for any `DataSendPoint` that references a provided `DataElementPrototype`

with `isQueued = false` that is marked as invalidatable and `canInvalidate` is enabled.

[rte_sws_ext_2682] The `Rte_Invalidate` API may only be used by the runnable that contains the corresponding `DataSendPoint`

Description: The `Rte_Invalidate` API takes no parameters other than the instance handle – the return value is used to indicate the success, or otherwise, of the API call to the caller.

[rte_sws_1231] When COM is used for communication and the `DataElementPrototype` is primitive the COM API function `Com_InvalidateSignal` shall be called for invalidation.

[rte_sws_5063] When COM is used for communication and the `DataElementPrototype` is composite the COM API function `Com_InvalidateShadowSignal` shall be called iteratively for invalidation of all the primitive parts of the composite `DataElementPrototype`.

The behavior required when COM is not used for communication is described in Section 4.3.1.8.

Return Value: The return value is used to indicate the “OK” status or errors detected by the RTE during execution of the `Rte_Invalidate` call.

- **[rte_sws_1207]** `RTE_E_OK` – No error occurred.
- **[rte_sws_1339]** `RTE_E_COM_STOPPED` – the RTE could not perform the operation because the COM service is currently not available (inter ECU communication only). RTE shall return `RTE_E_COM_STOPPED` when the corresponding COM service returns `COM_E_STOPPED`.

Notes: The API name includes an identifier `<p>_<o>` that is formed from the port and operation item names. See Section 5.2.6.4 for details on the naming convention.

The communication service configuration determines whether the signal receiver(s) receive an “invalid signal” notification or whether the invalidated signal is silently replaced by the signal’s initial value.

5.6.7 Rte_Feedback

Purpose: Provide access to acknowledgement notifications for explicit sender-receiver communication and to pass error notification to senders.

Signature: **[rte_sws_1083]**
`Std_ReturnType`
`Rte_Feedback_<p>_<o>([[IN Rte_Instance <instance>]])`

Where `<p>` is the port name and `<o>` the `DataElementPrototype` or `ModeDeclarationGroupPrototype` within the sender-receiver interface categorizing the port.

Existence: **[rte_sws_1283]** Acknowledgement is enabled for a provided `DataElementPrototype` or `ModeDeclarationGroupPrototype` by the presence of an `AcknowledgementRequest`.

[rte_sws_2676] Acknowledgement is enabled for a provided `ModeDeclarationGroupPrototype` when the `needsAck` attribute of the `ModeSwitchComSpec` is true.

[rte_sws_1284] A blocking `Rte.Feedback` API shall be generated for a provided `DataElementPrototype` if acknowledgement is enabled and a `WaitPoint` references a `DataSendCompletedEvent` that in turn references the `DataElementPrototype` or `ModeDeclarationGroupPrototype`.

[rte_sws_2677] A blocking `Rte.Feedback` API shall be generated for a provided `ModeDeclarationGroupPrototype` if acknowledgement is enabled and a `WaitPoint` references a `ModeSwitchedAckEvent` that in turn references the `ModeDeclarationGroupPrototype`.

[rte_sws_1285] A non-blocking `Rte.Feedback` API shall be generated for a provided `DataElementPrototype` if acknowledgement is enabled and a `DataSendPoint` references the `DataElementPrototype` but no `WaitPoint` references the `DataSendCompletedEvent` that references the `DataElementPrototype` or `ModeDeclarationGroupPrototype`.

[rte_sws_2678] A non-blocking `Rte.Feedback` API shall be generated for a provided `ModeDeclarationGroupPrototype` if acknowledgement is enabled and a `ModeSwitchPoint` references the `ModeDeclarationGroupPrototype` but no `ModeSwitchedAckEvent` references the `ModeDeclarationGroupPrototype`.

[rte_sws_1286] If acknowledgement is enabled for a provided `DataElementPrototype/ModeDeclarationGroupPrototype` and a `DataSendCompletedEvent` references a runnable entity as well as the `DataElementPrototype/ModeDeclarationGroupPrototype`, the runnable entity shall be activated when the transmission acknowledgement occurs or when a timeout was detected by the RTE. `rte_sws_1137`.

[rte_sws_2679] If acknowledgement is enabled for a provided `ModeDeclarationGroupPrototype` and a `ModeSwitchedAckEvent` references a runnable entity as well as the `ModeDeclarationGroupPrototype`, the runnable entity shall be activated when the mode switch acknowledgment occurs or when a timeout was detected by the RTE. `rte_sws_1137`.

Requirements `rte_sws_1286` and `rte_sws_2679` merely affect when the runnable is activated – an API call should still be created, according to requirements `rte_sws_1285` and `rte_sws_2678` to actually read the data.

[rte_sws_1287] A `DataSendCompletedEvent` or `ModeSwitchedAckEvent` that references a runnable entity and is referenced by a `WaitPoint` shall be an invalid configuration.

[rte_sws_ext_2687] A blocking `Rte_Feedback` API may only be used by the runnable that contains the corresponding `WaitPoint`

Description: The `Rte_Feedback` API takes no parameters other than the instance handle – the return value is used to indicate the acknowledgement status to the caller.

The `Rte_Feedback` API applies only to explicit sender-receiver communication.

Return Value: The return value is used to indicate the “status” status and errors detected by the RTE during execution of the `Rte_Feedback` call.

- **[rte_sws_1084]** `RTE_E_NO_DATA` – (non-blocking read) no data returned and no other error occurred when the feedback read was attempted.
- **[rte_sws_3774]** `RTE_E_COM_STOPPED` – (inter-ECU only) no data was returned within the specified timeout because the corresponding IPDU group was disabled.
- **[rte_sws_1086]** `RTE_E_TRANSMIT_ACK` – A transmission or mode switched acknowledgment has been received from the communication service. For intra-ECU communication this value is always returned even if a queue overflow occurred.

For intra ECU communication of mode switches, this indicates, that the runnables on the transition have been executed and the mode disablings have been switched to the new mode (see `rte_sws_2587`).

The `RTE_E_TRANSMIT_ACK` return value is not considered to be an error but rather indicates correct operation of the API call.

When `RTE_E_NO_DATA` occurs, a component is free to reinvoke `Rte_Feedback` and thus repeat the attempt to read the feedback status.

Notes: The API name includes an identifier `<p>_<o>` that indicates the read access point name and is formed from the port and operation item names. See Section 5.2.6.4 for details on the naming convention.

If multiple transmissions on the same port/element are outstanding it is not possible to determine which is acknowledged first. If this is important, transmissions should be serialized with the next occurring only when the previous transmission has been acknowledged or has timed out.

5.6.8 Rte_Read

Purpose: Performs an “explicit” read on a sender-receiver communication data element with “data” semantics (isQueued = false).

Signature: **[rte_sws_1091]**
`Std_ReturnType
Rte_Read_<p>_<o>([IN Rte_Instance <instance>],
OUT <data>)`

Where <p> is the port name and <o> the data element within the sender-receiver interface categorizing the port.

Existence: **[rte_sws_1289]** A non-blocking `Rte_Read` API shall be generated if a `DataReceivePoint` references a required `DataElementPrototype` with ‘data’ semantics (isQueued = false).

[rte_sws_1291] A `WaitPoint` that references a `DataReceivedEvent` that in turn references a required `DataElementPrototype` with ‘data’ semantics (isQueued = false) shall be considered an invalid configuration.

[rte_sws_1292] When a `DataReceivedEvent` references a `RunnableEntity` and a required `DataElementPrototype` and no `WaitPoint` references the `DataReceivedEvent`, the runnable entity shall be activated when the data is received. `rte_sws_1135`.

[rte_sws_ext_2683] The `Rte_Read` API may only be used by the runnable that contains the corresponding `DataReceivePoint`

Requirement `rte_sws_1292` merely affects when the runnable is activated – an API call should still be created, according to requirement `rte_sws_1288` or `rte_sws_1289` as appropriate, to actually read the data.

[rte_sws_1313] A `DataReceivedEvent` that references a runnable entity and is referenced by a `WaitPoint` shall be an invalid configuration.

Description: The `Rte_Read` API call includes exactly one OUT parameter to pass back the received data. The pointer to the OUT parameter must remain valid until the API call returns.

Return Value: The return value is used to indicate errors detected by the RTE during execution of the `Rte_Read` or `Rte_Receive` API call or errors detected by the communication system.

- **[rte_sws_1093]** `RTE_E_OK` – data read successfully.
- **[rte_sws_2626]** `RTE_E_INVALID` – data element invalid.
- **[rte_sws_2703]** `RTE_E_MAX_AGE_EXCEEDED` – data element outdated. This Overlaid Error can be combined with any of the above error codes.

Notes: The API name includes an identifier `<p>_<o>` that indicates the read access point name and is formed from the port and operation item names. See Section 5.2.6.4 for details on the naming convention.

5.6.9 Rte_Receive

Purpose: Performs an “explicit” read on a sender-receiver communication data element with “event” semantics (`isQueued = true`).

[rte_sws_1092]

```
Std_ReturnType
Rte_Receive_<p>_<o>([IN Rte_Instance <instance>],
                  OUT <data>)
```

Where `<p>` is the port name and `<o>` the data element within the sender-receiver interface categorizing the port.

Existence: **[rte_sws_1288]** A non-blocking `Rte_Receive` API shall be generated if a `DataReceivePoint` references a required `DataElementPrototype` with ‘event’ semantics (`isQueued = true`).

[rte_sws_1290] A blocking `Rte_Receive` API shall be generated if a `DataReceivePoint` references a required `DataElementPrototype` with ‘event’ semantics (`isQueued = true`) that is, in turn, referenced by a `DataReceivedEvent` and the `DataReceivedEvent` is referenced by a `WaitPoint`.

When a `DataReceivedEvent` references a `RunnableEntity` and a required `DataElementPrototype` and no `WaitPoint` references the `DataReceivedEvent`, the runnable entity shall be activated when the event is received. `rte_sws_1292` `rte_sws_1135`.

Requirement `rte_sws_1292` merely affects when the runnable is activated – an API call should still be created, according to requirement `rte_sws_1288` or `rte_sws_1289` as appropriate, to actually read the data.

[rte_sws_ext_2684] The `Rte_Receive` API may only be used by the runnable that contains the corresponding `DataReceivePoint`

A `DataReceivedEvent` that references a runnable entity and is referenced by a `WaitPoint` shall be an invalid configuration. `rte_sws_1313`

Description: The `Rte_Receive` API call includes exactly one OUT parameter to pass back the received data.

The pointer to the OUT parameter must remain valid until the API call returns.

Return Value: The return value is used to indicate errors detected by the RTE during execution of the `Rte_Receive` API call or errors detected by the communication system.

- **[rte_sws_2598]** `RTE_E_OK` – data read successfully.
- **[rte_sws_1094]** `RTE_E_NO_DATA` – (explicit non-blocking read) no data returned and no other error occurred when the read was attempted.
- **[rte_sws_1095]** `RTE_E_TIMEOUT` – (explicit blocking read) no data returned and no other error occurred when the read was attempted.
- **[rte_sws_2572]** `RTE_E_LOST_DATA` – Indicates that some incoming data has been lost due to an overflow of the receive queue or due to an error of the underlying communication layers. This is not an error of the data returned in the parameters. This `Overlaid Error` can be combined with any of the above.

The `RTE_E_NO_DATA` and `RTE_E_TIMEOUT` return value are not considered to be errors but rather indicate correct operation of the API call.

Notes: The API name includes an identifier `<p>_<o>` that indicates the read access point name and is formed from the port and operation item names. See Section 5.2.6.4 for details on the naming convention.

5.6.10 Rte_Call

Purpose: Initiate a client-server communication.

Signature: **[rte_sws_1102]**

```
Std_ReturnType
Rte_Call_<p>_<o>([IN Rte_Instance <instance>],
                [IN|IN/OUT|OUT] <data_1>...
                [IN|IN/OUT|OUT] <data_n>)
```


Where <p> is the port name and <o> the operation within the client-server interface categorizing the port.

Existence: [rte_sws_1293] A synchronous `Rte_Call` API shall be generated if a `SynchronousServerCallPoint` references a required `OperationPrototype`.

[rte_sws_1294] An asynchronous `Rte_Call` API shall be generated if an `AsynchronousServerCallPoint` references a required `OperationPrototype`.

A configuration that includes both synchronous and asynchronous `ServerCallPoints` for a given `OperationPrototype` is invalid (rte_sws_3014).

[rte_sws_ext_2685] The `Rte_Call` API may only be used by the runnable that contains the corresponding `ServerCallPoint`

Description: Client function to initiate client-server communication. The `Rte_Call` API is used for both synchronous and asynchronous calls.

The `Rte_Call` API includes zero or more IN, IN/OUT and OUT parameters. IN parameters are passed by value for primitive data types and by reference for all other types, OUT parameters are always by reference and IN/OUT parameters are passed by value when they are primitive data types and the call is asynchronous and by reference for all other cases.

The pointers to all parameters passed by reference must remain valid until the API call returns.

Return Value: [rte_sws_1103] The return value shall be used to indicate infrastructure errors detected by the RTE during execution of the `Rte_Call` call and, for synchronous communication, infrastructure and application errors during execution of the server.

- [rte_sws_1104] `RTE_E_OK` – The API call completed successfully.
- [rte_sws_1105] `RTE_E_LIMIT` – The client has multiple outstanding asynchronous client-server invocations in the same server call point. The server invocation shall be discarded, the buffers of the return parameters shall not be modified (see also rte_sws_2658).
- [rte_sws_1106] `RTE_E_COM_STOPPED` – the RTE could not perform the operation because the COM service is currently not available (inter ECU communication only). RTE shall return `RTE_E_COM_STOPPED` when the corresponding COM service returns `COM_E_STOPPED`. The buffers of the return parameters shall not be modified.

- **[rte_sws_1107]** `RTE_E_TIMEOUT` – (synchronous inter-task and inter-ECU only) No reply was received within the configured timeout. The buffers of the return parameters shall not be modified.
- **[rte_sws_2577]** The application error (synchronous client-server) from a server shall only be returned if none of the above infrastructure errors (other than `RTE_E_OK`) have occurred.

Note that the `RTE_E_OK` return value indicates that the `Rte_Call` API call completed successfully. In case of a synchronous client server call it also indicates successful processing of the request by the server.

An asynchronous server invocation is considered to be outstanding until either the client retrieved the result successfully, a timeout was detected by the RTE in inter-ECU communication or the server runnable has terminated after a timeout was detected in intra-ECU communication.

When the `RTE_E_TIMEOUT` error occurs, RTE shall discard any subsequent responses to that request, (see `rte_sws_2657`).

Notes: **[rte_sws_1109]** The interface operation's OUT parameters shall be omitted for an *asynchronous* call.

For asynchronous communication the `Rte_Call` should include only IN and IN/OUT parameters – the OUT parameters are required when the client collects the result (e.g. using `Rte_Result`).

5.6.11 Rte_Result

Purpose: Get the result of an asynchronous client-server call.

Signature: **[rte_sws_1111]**
`Std_ReturnType`
`Rte_Result_<p>_<o>([IN Rte_Instance <instance>],`
`[OUT <param 1>]...`
`[OUT <param n>])`

Where `<p>` is the port name and `<o>` the operation within the client-server interface categorizing the port.

The signature can include zero or more OUT parameters depending on the signature of the operation in the client-server interface.

Existence: **[rte_sws_1296]** A non-blocking `Rte_Result` API shall be generated if an `AsynchronousServerCallReturnsEvent` references a required `OperationPrototype` and no `WaitPoint` references the `AsynchronousServerCallReturnsEvent`.

[rte_sws_1297] A blocking `Rte_Result` API shall be generated if an `AsynchronousServerCallReturnsEvent` references a required `OperationPrototype` and a `WaitPoint` references the `AsynchronousServerCallReturnsEvent`.

[rte_sws_ext_2686] The blocking `Rte_Result` API may only be used by the runnable that contains the corresponding `WaitPoint`

[rte_sws_1298] If an `AsynchronousServerCallReturnsEvent` references a `RunnableEntity` and a required `OperationPrototype` the runnable entity shall be activated when the operation's result is available or when a timeout was detected by the RTE `rte_sws_1133`.

Requirement `rte_sws_1298` merely affects when the runnable is activated – an API call should still be created to actually read the reply based on requirement `rte_sws_1296`.

[rte_sws_1312] An `AsynchronousServerCallReturnsEvent` that references a runnable entity and is referenced by a `WaitPoint` is invalid.

Description: The `Rte_Result` API is used by a client to collect the result of an *asynchronous* client-server communication.

The `Rte_Result` API includes zero or more OUT parameters to pass back results.

The pointers to all parameters passed by reference must remain valid until the API call returns.

Return Value: The return value is used to indicate errors from either the `Rte_Result` call itself or communication errors detected before the API call was made.

- **[rte_sws_1112]** `RTE_E_OK` – The API call completed successfully.
- **[rte_sws_1113]** `RTE_E_NO_DATA` – (non-blocking read) The server's result is not available but no other error occurred within the API call. The buffers for the OUT parameters shall not be modified.
- **[rte_sws_1114]** `RTE_E_TIMEOUT` – The server's result is not available within the specified timeout but no other error occurred within the API call. The buffers for the OUT parameters shall not be modified.
- **[rte_sws_3606]** `RTE_E_COM_STOPPED` – the RTE could not perform the operation because the COM service is currently not available (inter ECU communication only). RTE shall return `RTE_E_COM_STOPPED` when the corresponding COM service returns `COM_E_STOPPED`. The server's result has *not* been suc-

cessfully retrieved from the communication service. The buffers of the return parameters shall not be modified.

- **[rte_sws_2578]** Application Errors – The error code of the server shall only be returned, if none of the above infrastructure errors or indications have occurred.

The `RTE_E_NO_DATA` and `RTE_E_TIMEOUT` return value are not considered to be errors but rather indicate correct operation of the API call.

When the `RTE_E_TIMEOUT` error occurs, RTE shall discard any subsequent responses to that request, (see `rte_sws_2657`).

When `RTE_E_NO_DATA` occurs, a component is free to invoke `Rte_Result` again and thus repeat the attempt to read the server's result.

Notes: The API name includes an identifier `<p>_<o>` that indicates the read access point name and is formed from the port and operation item names. See Section 5.2.6.4 for details on the naming convention.

5.6.12 Rte_Pim

Purpose: Provide access to the defined per-instance memory (section) of a software component.

Signature: **[rte_sws_1118]**
`<type>`
`Rte_Pim_<name>([IN Rte_Instance <instance>])`

Where `<name>` is the (short) name of the per-instance name.

Existence: **[rte_sws_1299]** An `Rte_PIM` API shall be created for each defined `PerInstanceMemory` within the AUTOSAR software-component (description).

Description: The `Rte_PIM` API provides access to the per-instance memory (section) defined in the context of a `InternalBehavior` of a software-component description.

Return Value: **[rte_sws_1119]** The API returns a typed reference (in C a typed pointer) to the per-instance memory.

Notes: The software-component shall define the return type `<type>` in the attribute `<typeDefinition>` of `PerInstanceMemory`, if it is a complex AUTOSAR data type. It is assumed that this attribute contains a String that represents a C type definition (typedef) in valid C syntax (see `rte_sws_2304`).

5.6.13 Rte_CData

Purpose: Provide access to the calibration parameter an AUTOSAR software-component defined internally. The CalprmElementPrototype is used to define software component internal calibration parameters. Internal because the CalprmElementPrototype cannot be reused outside the software-component. Access is read-only. It can be configured for each calibration parameter individually if it is shared by all instances of an AUTOSAR software-component or if each instance has an own data value associated with it.

Signature: **[rte_sws_1252]**
`<return>`
`Rte_CData_<name>([[IN Rte_Instance <instance>]])`

Where `<name>` is the calibration parameter name.

Existence: **[rte_sws_1300]** An `Rte_CData` API shall be created for each defined `CalprmElementPrototype` within an AUTOSAR software-component.

Description: The `Rte_CData` API provides access to the defined calibration parameter within a software-component. The actual data values for a software-component instance may be set after component compilation.

Return Value: **[rte_sws_1254]** The `Rte_CData` API shall return access to the calibration parameter value. For the type of the return value refer to `rte_sws_5195`, `rte_sws_5198`, `rte_sws_5207`, `rte_sws_5208`, `rte_sws_5209`.

[rte_sws_3927] If the attribute "perInstanceCalprm" of a `CalprmElementPrototype` of a software-component is set to "FALSE", the return value shall provide access to one common calibration parameter for all instances.

[rte_sws_3952] If the attribute "perInstanceCalprm" of a `CalprmElementPrototype` of a software-component is set to "TRUE", the return value of the `Rte_CaData` API shall provide access to the instance specific calibration parameter.

Notes: None.

5.6.14 Rte_Calprm

Purpose: Provide access to the calibration parameters defined by an AUTOSAR `CalprmComponentType`. Access is read-only.

Signature: **[rte_sws_3928]**

```
<return>
Rte_Calprm_<p>_<name>([IN Rte_Instance <instance>])
```

Where <p> is the port name and <name> is the calibration parameter name.

Existence: [rte_sws_3929] An `Rte_Calprm` API shall be created for each defined `CalprmElementPrototype` within an `AUTOSAR CalprmComponentType`.

Description: The `Rte_Calprm` API provides access to the defined calibration parameter within a `CalprmComponentType`. The actual data values for a `CalprmComponentType` instance may be set after `CalprmComponentType` compilation.

Return Value: [rte_sws_3930] The `Rte_Calprm` API shall return access to the calibration parameter value. For the type of the return value refer to `rte_sws_5195`, `rte_sws_5198`, `rte_sws_5207`, `rte_sws_5208`, `rte_sws_5209`.

Notes: None.

5.6.15 Rte_IRead

Purpose: Provide **read** access to the data elements defined with `DataReadAccess` semantics.

Signature: [rte_sws_3741]

```
<return>
Rte_IRead_<re>_<p>_<d>([IN Rte_Instance])
```

Where <re> is the runnable entity name, <p> the port name and <d> the data element name.

Existence: [rte_sws_1301] An `Rte_IRead` API shall be created for a required `DataElementPrototype` if the `RunnableEntity` has `DataReadAccess` that refers to the `DataElementPrototype`.

Description: The `Rte_IRead` API provides access to the data elements declared as accessed by a runnable using `DataReadAccess`. The API function is guaranteed to have constant execution time and therefore can also be used within category 1A runnable entities.

No error information is provided by this API. If required, the error status can be picked up with a separate API, see 5.6.19

The data value can always be read. To provide the required consistency the API provides access to a *copy* of the data data element for which it's guaranteed that it never changes during the actual execution of the runnable entity.

Implicit data read access by a SW-C should always return defined data.

[rte_sws_1268] The RTE shall ensure that implicit read accesses will not deliver undefined data item values.

In case where there may be an implicit read access before the first data reception an initial value has to be provided as the result of this implicit read access.

Return Value: **[rte_sws_3593]** The `Rte_IRead` API shall return access to the value of the data element. For the type of the return value refer to `rte_sws_5195`, `rte_sws_5198`, `rte_sws_5207`, `rte_sws_5208`, `rte_sws_5209`.

Notes: None.

5.6.16 Rte_IWrite

Purpose: Provide **write** access to the data elements defined with `DataWriteAccess` semantics.

Signature: **[rte_sws_3744]**

```
void
Rte_IWrite_<re>_<p>_<d>([IN RTE_Instance],
                        IN <type>)
```

Where `<re>` is the runnable entity name, `<p>` the port name and `<d>` the data element name. If the datatype of the data element is a primitive type, `<type>` is the datatype of the data element. If the datatype is a composite datatype, `<type>` is the pointertype of the data element's datatype.

Existence: **[rte_sws_1302]** An `Rte_IWrite` API shall be created for a provided `DataElementPrototype` if the `RunnableEntity` has `DataWriteAccess` that refers to the `DataElementPrototype`.

Description: The `Rte_IWrite` API provides write access to the data elements declared as accessed by a runnable using `DataWriteAccess`. The API function is guaranteed to be have constant execution time and therefore can also be used within category 1A runnable entities.

No access error information is required for the user – the value can always be written. To provide the required write-back semantics the RTE only makes written values available to other entities after the writing runnable entity has terminated.

[rte_sws_3746] The `Rte_IWrite` API call include exactly one IN parameter for the data element – this is passed by value for primitive data types and by reference for all other types.

Return Value: [rte_sws_3747] `Rte_IWrite` has no return value.

For C/C++ `rte_sws_3747` means using a return type of `void`.

Notes: None.

5.6.17 Rte_IWriteRef

Purpose: Provide a reference to the data elements defined with `DataWriteAccess` semantics.

Signature: [rte_sws_5509]

<return>

`Rte_IWriteRef_<re>_<p>_<d>([IN RTE_Instance])`

Where <re> is the runnable entity name, <p> the port name and <d> the data element name.

Existence: [rte_sws_5510] An `Rte_IWriteRef` API shall be created for a provided `DataElementPrototype` if the `RunnableEntity` has `DataWriteAccess` that refers to the `DataElementPrototype`.

Description: The `Rte_IWriteRef` API returns a reference to the data elements declared as accessed by a runnable using `DataWriteAccess`. The reference can be used by the runnable to directly update the corresponding data elements. This is especially useful for data elements of complex types or strings. The API function is guaranteed to be have constant execution time and therefore can also be used within category 1A runnable entities.

No error information is required for the user. To provide the required write-back semantics the RTE only makes written values available to other entities after the writing runnable entity has terminated.

Return Value: [rte_sws_5511] `Rte_IWriteRef` returns a reference to the corresponding data element.

[rte_sws_5512] The return type of `Rte_IWriteRef` is dependent on the data element type. For a primitive data type (with the exception of a string) it is a pointer to the data element type. For composite data types refer to `rte_sws_5195`, `rte_sws_5198`, `rte_sws_5207`, `rte_sws_5208`, `rte_sws_5209`.

Notes: None.

5.6.18 Rte_IInvalidate

- Purpose:** Invalidate a data element defined with `DataWriteAccess` semantics.
- Signature:** **[rte_sws_3800]**
`void`
`Rte_IInvalidate_<re>_<p>_<d>([IN Rte_Instance <instance>])`
 Where `<re>` is the runnable entity name, `<p>` the port name and `<d>` the data element name.
- Existence:** **[rte_sws_3801]** An `Rte_IInvalidate` API shall be created for a provided `DataElementPrototype` if the `RunnableEntity` has `DataWriteAccess` that refers to the `DataElementPrototype` and `canInvalidate` is enabled.
- Description:** The `Rte_IInvalidate` API takes no parameters other than the instance handle – the return value is used to indicate the success, or otherwise, of the API call to the caller.
- [rte_sws_3802]** In case of a primitive `DataElementPrototype` the `Rte_IInvalidate` shall be implemented as a macro that writes the invalid value `rte_sws_in_5031` to the buffer.
- [rte_sws_5064]** In case of a composite `DataElementPrototype` the `Rte_IInvalidate` shall be implemented as a macro that writes the invalid values `rte_sws_in_5031` of every primitive part of the composition to the buffer.
- [rte_sws_3778]** If `Rte_IInvalidate` is followed by an `Rte_IWrite` for the same data element prototype call or vice versa, the RTE shall use the last value written before the runnable entity terminates (last-is-best semantics).
- `rte_sws_3778` states that an `Rte_IWrite` overrules an `Rte_IInvalidate` call if it occurs after the `Rte_IInvalidate`, since `Rte_IWrite` overwrites the contents of the internal buffer for the data element prototype before it is made known to other runnable entities.
- Return Value:** **[rte_sws_3803]** `Rte_IInvalidate` has no return value.
- For C/C++ `rte_sws_3803` means using a return type of `void`.
- Notes:** The communication service configuration determines whether the signal receiver(s) receive an “invalid signal” notification or whether the invalidated signal is silently replaced by the signal’s initial value.

5.6.19 Rte_IStatus

Purpose: Provide the error status of a data element defined with `DataReadAccess` semantics.

Signature: **[rte_sws_2599]**

`Std_ReturnType`

`Rte_IStatus_<re>_<p>_<d> ([IN Rte_Instance])`

Where `<re>` is the runnable entity name, `<p>` the port name and `<d>` the data element name.

Existence: **[rte_sws_2600]** An `Rte_IStatus` API shall be created for a required `DataElementPrototype` if a `RunnableEntity` has `DataReadAccess` referring to the `DataElementPrototype` and if either

- data element outdated notification or
- data element invalidation

is activated for this data element.

[rte_sws_ext_2601] The `Rte_IStatus` API shall only be used by a `RunnableEntity` that either has a `DataReadAccess` referring to the `DataElementPrototype` or is triggered by a `DataReceiveErrorEvent` referring to the `DataElementPrototype`.

Description: The `Rte_IStatus` API provides access to the current status of the data elements declared as accessed by a runnable using `DataReadAccess`. The API function is guaranteed to have constant execution time and therefore can also be used within category 1A runnable entities.

To provide the required consistency access by a runnable is to a *copy* of the status together with the data that is guaranteed never to be modified by the RTE during the lifetime of the runnable entity.

Return Value: The return value is used to indicate errors detected by the communication system.

- **[rte_sws_2602]** `RTE_E_OK` – no errors.
- **[rte_sws_2603]** `RTE_E_INVALID` – data element invalid.
- **[rte_sws_2604]** `RTE_E_MAX_AGE_EXCEEDED` – data element outdated. This Overlaid Error can be combined with any of the above error codes.

Notes: None.

5.6.20 Rte_IrvIRead

Purpose: Provide **read** access to the *InterRunnableVariables with implicit* behavior of an AUTOSAR SW-C.

Signature: **[rte_sws_3550]**

<return>

Rte_IrvIRead_<re>_<name> ([IN RTE_Instance <instance>])

Where <re> is the name of the runnable entity the API might be used in, <name> is the name of the InterRunnableVariables.

Existence: **[rte_sws_1303]** An Rte_IrvIRead API shall be created for each read InterRunnableVariable.

Description: The Rte_IrvIRead API provides read access to the defined InterRunnableVariables with *implicit* behavior within a component description.

The return value is used to deliver the requested data value. The return value is not required to pass error information to the user because no inter-ECU communication is involved and there will always be a readable value present.

Requirement rte_sws_3581 is valid for InterRunnableVariables with implicit and InterRunnableVariables with explicit behavior:

[rte_sws_3581] The RTE has to ensure that read accesses to an InterRunnableVariables won't deliver undefined data item values. In case write access before read access cannot be guaranteed by configuration an initial values for the InterRunnableVariable has to be written to it.

This initial value has to be an input for the RTE generator and might be initially defined in the AUTOSAR SW-C description.

Return Value: **[rte_sws_3552]** The Rte_IrvIRead call returns the actual value of the accessed InterRunnableVariable.

The return type of Rte_IrvIRead is dependent on the InterRunnableVariable data type. Thus the component does not need to use type casting to convert access the InterRunnableVariable data.

[rte_sws_3556] The return value of the Rte_IrvIRead API call shall pass a value.

[rte_sws_3558] The Rte_IrvIRead API call does not support complex data types.

Notes: The runnable entity name in the signature allows runnable context specific optimizations.

The concept of `InterRunnableVariables` is explained in section 4.2.4.6. More details about `InterRunnableVariables` with *implicit* behavior is explained in section 4.2.4.6.1.

5.6.21 Rte_IrvIWrite

Purpose: Provide **write** access to the *InterRunnableVariables* with *implicit* behavior of an AUTOSAR SW-C.

Signature: **[rte_sws_3553]**

```
void
Rte_IrvIWrite_<re>_<name>([IN RTE_Instance <instance>],
                          IN <data>)
```

Where `<re>` is the name of the runnable entity the API might be used in, `<name>` is the name of the `InterRunnableVariable` to access and `<data>` is the placeholder for the data the `InterRunnableVariable` shall be set to.

Existence: **[rte_sws_1304]** An `Rte_IrvIWrite` API shall be created for each written `InterRunnableVariable`.

Description: The `Rte_IrvIWrite` API provides write access to the `InterRunnableVariables` with *implicit* behavior within a component description. The runnable entity name in the signature allows runnable context specific optimizations.

The data given by `Rte_IrvIWrite` is dependent on the `InterRunnableVariable` data type. Thus the component does not need to use type casting to write the `InterRunnableVariable`.

The return value is unused. The return value is not required to pass error information to the user because no inter-ECU communication is involved and the value can always be written.

[rte_sws_3557] The `Rte_IrvIWrite` API call include exactly one IN parameter for the data element - which is a pass by value.

[rte_sws_3559] The `Rte_IrvIWrite` API call does not support complex data types.

Return Value: **[rte_sws_3555]** `Rte_IrvIWrite` shall have no return value.

For C/C++, requirement `rte_sws_3555` means using a return type of `void`.

Notes: The runnable entity name in the signature allows runnable context specific optimizations.

The concept of InterRunnableVariables is explained in section 4.2.4.6. Further details about InterRunnableVariables with *implicit* behavior are explained in Section 4.2.4.6.1.

5.6.22 Rte_IrvRead

Purpose: Provide **read** access to the *InterRunnableVariables with explicit behavior* of an AUTOSAR SW-C.

Signature: **[rte_sws_3560]**

```
<return>
Rte_IrvRead_<re>_<name> ([IN RTE_Instance <instance>])
```

Where <re> is the name of the runnable entity the API might be used in, <name> is the name of the InterRunnableVariables.

Existence: **[rte_sws_1305]** An `Rte_IrvIRead` API shall be created for each read InterRunnableVariable using explicit access.

Description: The `Rte_IrvRead` API provides read access to the defined InterRunnableVariables with *explicit* behavior within a component description.

The return value is used to deliver the requested data value. The return value is not required to pass error information to the user because no inter-ECU communication is involved and there will always be a readable value present.

Return Value: **[rte_sws_3562]** The `Rte_IrvRead` call returns the actual value of the accessed InterRunnableVariable.

The return type of `Rte_IrvRead` is dependent on the InterRunnableVariable data type. Thus the component does not need to use type casting to convert access the InterRunnableVariable data.

[rte_sws_3563] The return value of the `Rte_IrvRead` API call shall pass a value.

[rte_sws_3564] The `Rte_IrvRead` API call does not support complex data types.

Notes: The runnable entity name in the signature allows runnable context specific optimizations.

The concept of InterRunnableVariables is explained in section 4.2.4.6. Further details about InterRunnableVariables with *explicit* behavior are explained in Section 4.2.4.6.2.

5.6.23 Rte_IrvWrite

Purpose: Provide **write** access to the *InterRunnableVariables with explicit behavior* of an AUTOSAR SW-C.

Signature: **[rte_sws_3565]**

```
void
Rte_IrvWrite_<re>_<name>([IN RTE_Instance <instance>],
                        IN <data>)
```

Where <re> is the name of the runnable entity the API might be used in, <name> is the name of the InterRunnableVariable to access and <data> is the placeholder for the data the InterRunnableVariable shall be set to.

Existence: **[rte_sws_1306]** An `Rte_IrvIWrite` API shall be created for each written InterRunnableVariable using explicit access.

Description: The `Rte_IrvWrite` API provides write access to the InterRunnableVariables with *explicit* behavior within a component description.

The data given by `Rte_IrvWrite` is dependent on the InterRunnableVariable data type. Thus the component does not need to use type casting to write the InterRunnableVariable.

The return value is unused. The return value is not required to pass error information to the user because no inter-ECU communication is involved and the value can always be written.

[rte_sws_3567] The `Rte_IrvWrite` API call include exactly one IN parameter for the data element - which is a pass by value.

[rte_sws_3568] The `Rte_IrvWrite` API call does not support complex data types.

Return Value: **[rte_sws_3569]** `Rte_IrvWrite` shall have no return value.

For C/C++, requirement `rte_sws_3569` means using a return type of `void`.

Notes: The runnable entity name in the signature allows runnable context specific optimizations.

The concept of InterRunnableVariables is explained in section 4.2.4.6. Further details about InterRunnableVariables with *explicit* behavior are explained in Section 4.2.4.6.2.

5.6.24 Rte_Enter

Purpose: Enter an exclusive area.

- Signature:** **[rte_sws_1120]**
`void
Rte_Enter_<name>([IN Rte_Instance <instance>])`
Where <name> is the exclusive area name.
- Existence:** **[rte_sws_1307]** An `Rte_Enter` API shall be created for each `ExclusiveArea` that is declared `RunnableEntityCanEnterExclusiveArea`.
- Description:** The `Rte_Enter` API call is invoked by an AUTOSAR software-component to define the start of an exclusive area.
- Return Value:** None.
- Notes:** The RTE is not required to support nested invocations of `Rte_Enter` for the same exclusive area.
- [rte_sws_1122]** The RTE shall permit calls to `Rte_Enter` and `Rte_Exit` to be nested as long as regions are exited in the reverse order they were entered.
- Within the AUTOSAR OS an attempt to lock a resource cannot fail because the lock is already held. The lock attempt can only fail due to configuration errors (e.g. caller not declared as accessing the resource) or invalid handle. Therefore the return type from this function is `void`.

5.6.25 Rte_Exit

- Purpose:** Leave an exclusive area.
- Signature:** **[rte_sws_1123]**
`void
Rte_Exit_<name>([IN Rte_Instance <instance>])`
Where <name> is the exclusive area name.
- Existence:** **[rte_sws_1308]** An `Rte_Exit` API shall be created for each `ExclusiveArea` that is declared `RunnableEntityCanEnterExclusiveArea`.
- Description:** The `Rte_Exit` API call is invoked by an AUTOSAR software-component to define the end of an exclusive area.
- Return Value:** None.
- Notes:** The RTE is not required to support nested invocations of `Rte_Exit` for the same exclusive area.
- Requirement `rte_sws_1122` permits calls to `Rte_Enter` and `Rte_Exit` to be nested as long as regions are exited in the reverse order they were entered.

5.6.26 Rte_Mode

Purpose: Provides the currently active mode of a mode port.

Signature: [rte_sws_2628]

```
Rte_ModeType_<m>  
Rte_Mode_<p>_<o>([[IN Rte_Instance <instance>]])
```

Where <m> is the ModeDeclarationGroup name, <p> is the port name, and <o> the ModeDeclarationGroupPrototype name within the sender-receiver interface categorizing the port.

Existence: [rte_sws_2629] An `Rte_Mode` API shall be created for each required ModeDeclarationGroupPrototype and for each provided ModeDeclarationGroupPrototype.

Description: The `Rte_Mode` API tells the AUTOSAR Software-Component which mode of a ModeDeclarationGroup of a given port is currently active. This is the information that the RTE uses for the ModeDisablingDependencies. A new mode will not be indicated immediately after the reception of a mode switch notification from a mode manager, see section 4.4.4. During mode transitions, i.e. during the execution of runnables that are triggered on exiting one mode or on entering the next mode, overlapping mode disablings of two modes are active. In this case, the `Rte_Mode` will return `RTE_TRANSITION_<ModeDeclarationGroup>`.

The `Rte_Mode` will return the same mode for all mode ports that are connected to the same mode port of the mode manager (see `rte_sws_2630`).

Return Value: [rte_sws_2660] The `Rte_Mode` API shall return the following values:

- during mode transitions:
`RTE_TRANSITION_<ModeDeclarationGroup>`,
where <ModeDeclarationGroup> is the short name of the ModeDeclarationGroup.
- else:
`RTE_MODE_<ModeDeclarationGroup>_<ModeDeclaration>`,
where <ModeDeclarationGroup> is the short name of the ModeDeclarationGroup and <ModeDeclaration> is the short name of the currently active ModeDeclaration

currently active mode of the given instance of a ModeDeclarationGroupPrototype.

Notes: None.

5.7 Runnable Entity Reference

An AUTOSAR component defines one or more “runnable entities”. A runnable entity is a piece of code with a single entry point and an associate set of data. A software-component description provides definitions for each runnable entity within the software-component.

For components implemented using C or C++ the entry point of a runnable entity is implemented by a function with global scope defined within a software-component’s source code. The following sections consider the function signature and prototype.

5.7.1 Signature

The definition of all runnable entities, whatever the `RTEEvent` that triggers their execution, follows the same basic form.

[rte_sws_1126]

```
<void|Std_ReturnType> <name>([IN Rte_Instance <instance>],  
                             [role parameters])
```

Where `<name>`³ is the symbol describing the runnable’s entry point `rte_sws_in_0053`. The definition of the *role parameters* is defined in Section 5.7.3.

Section 5.2.6.4 contains details on a recommended naming conventions for runnable entities based on the `RTEEvent` that triggers the runnable entity. The recommended naming convention makes explicit the functions that implement runnable entities as well as clearly associating the runnable entity and the applicable data element or operation.

5.7.2 Entry Point Prototype

The RTE determines the required role parameters, and hence the prototype of the entry point, for a runnable entity based on information in the input information (see Appendix B). The entry point defined in the component source *must* be compatible with the parameters passed by the RTE when the runnable entity is triggered by the RTE and therefore the RTE generator is required to emit a prototype for the function.

[rte_sws_1132] The RTE generator shall emit a prototype for the runnable entity’s entry point in the application header file.

The prototype for a function implementing the entry point of a runnable entity is emitted for both “RTE Contract” and “RTE Generation” phases. The function name for the prototype is the runnable entity’s entry point. The prototype of the entry point function includes the runnable entity’s instance handle and its role parameters, see Figure 5.2.

³Runnable entities have two “names” associated with them in the Software-Component Template; the runnable’s identifier and the entry point’s symbol. The identifier is used to reference the runnable entity within the input data and the symbol used within code to identify the runnable’s implementation. In the context of a prototype for a runnable entity, “name” is the runnable entity’s entry point symbol.

[rte_sws_1016] The function implementing the entry point of a runnable entity shall define an instance handle as the first formal parameter.

The RTE will ensure that when the runnable entity is triggered the instance handle parameter indicates the correct component instance. The remaining parameters passed to the runnable entity depend on the `RTEEvent` that triggers execution of the runnable entity.

5.7.3 Role Parameters

The *role parameters* are optional and their presence and types depend on the `RTEEvent` that triggers the execution of the runnable entity. The role parameters that are necessary for each triggering `RTEEvent` are defined in Section 5.7.5.

5.7.4 Return Value

A function in C or C++ is required to have a return type. The RTE only uses the function return value to return application error codes of a server operation.

[rte_sws_1130] A function implementing a runnable entity entry point shall only have the return type `Std_ReturnType`, if the runnable entity represents a server operation and the AUTOSAR interface description of that client server communication lists potential application errors. All other functions implementing a runnable entity entry point shall have a return type of `void`.

[rte_sws_ext_2704] Only the least significant six bit of the return value of a server runnable shall be used by the application to indicate an error. The upper two bit shall be zero. See also `rte_sws_2573`.

5.7.5 Triggering Events

The RTE is the *sole* entity that can trigger the execution of a runnable entity. The RTE triggers runnable entities in response to different `RTEEvents`.

The most basic `RTEEvent` that can trigger a runnable entity is the `TimingEvent` that causes a runnable entity to be periodically triggered by the RTE. In contrast, the remaining `RTEEvents` that can trigger runnable entities all occur as a result of communication activity or as a result of mode switches.

The following subsections describe the conditions that can trigger execution of a runnable entity. For each triggering event the signature of the function (the “entry point”) that implements the runnable entity is defined. The signature definition includes two classes of parameters for each function;

1. The instance handle – the parameter type is always `Rte_Instance`. (rte_sws_1016)
2. The role parameters – used to pass information required by the runnable entity as a consequence of the triggering condition. The presence (and number) of role parameters depends solely on the triggering condition.

5.7.5.1 TimingEvent

Purpose: Trigger a runnable entity periodically at a rate defined within the software-component description.

Signature: **[rte_sws_1131]**
`void <name>([[IN Rte_Instance <instance>]])`

5.7.5.2 ModeSwitchEvent

Purpose: Trigger of a runnable entity as a result of a mode switch. See also sections 4.4.4 and 4.4.5 for reference.

Signature: **[rte_sws_2512]**
`void <name>([[IN Rte_Instance <instance>]])`

5.7.5.3 AsynchronousServerCallReturnsEvent

Purpose: Triggers a runnable entity used to “collect” the result and status information of an asynchronous client-server operation.

Signature: **[rte_sws_1133]**
`void <name>([[IN Rte_Instance <instance>]])`

Notes: The runnable entity triggered by an `AsynchronousServerCallReturnsEvent` `RTEEvent` should use the `Rte.Result` API to actually receive the result and the status of the server operation.

5.7.5.4 DataReceiveErrorEvent

Purpose: Triggers a runnable entity used to “collect” the error status of a `data element` with “data” semantics (`isQueued = false`) on the receiver side.

Signature: **[rte_sws_1359]**
`void <name>([[IN Rte_Instance <instance>]])`

Notes: The runnable entity triggered by a `DataReceiveErrorEvent` `RTEEvent` should use the `Rte_IStatus` API to actually read the status.

5.7.5.5 OperationInvokedEvent

Purpose: An `RTEEvent` that causes the RTE to trigger a runnable entity whose entry point provides an implementation for a client-server operation. This event occurs in response to a received request from a client to execute the operation.

Signature: **[rte_sws_1166]**

```
<void|Std_ReturnType> <name>
    ([IN Rte_Instance <instance>],
     [IN          <portDefArg 1>, ...
      IN          <portDefArg n>],
     [IN|INOUT|OUT] <param 1>, ...
     [IN|INOUT|OUT] <param n>)
```

Where `<portDefArg 1>, ..., <portDefArg n>` represent the port-defined argument values (see Section 4.3.2.4) and `<param 1>, ... <param n>` indicates the operation IN, INOUT and OUT parameters.

The data type of each port defined argument is taken from the software component template, as defined in `rte_sws_in_1361`.

Note that the port-defined argument values are optional, depending upon the server's internal behavior.

The operation parameters `<param 1>, ... <param n>` are the specified `ArgumentPrototypes` of the `OperationPrototype` that is associated with the `OperationInvokedEvent`. The operation parameters are ordered according to the `OperationPrototype`'s ordered list of the `ArgumentPrototypes`.

Return Value: If the AUTOSAR interface description of the client server communication lists possible error codes, these are returned by the function using the return type `Std_ReturnType`. If no error codes are defined for this interface, the return type shall be `void` (see `rte_sws_1130`).

This means that even if a runnable entity implementing a server "only" returns `E_OK`, application errors have to be defined. Else the return types do not match.

5.7.5.6 DataReceivedEvent

Purpose: A runnable entity triggered by the RTE to receive and process a signal received on a sender-receiver interface.

Signature: **[rte_sws_1135]**
`void <name>([[IN Rte_Instance <instance>]])`

Notes: The data or event is not passed as an additional parameter. Instead, the previously described reception API should be used to access the data/event. This approach permits the same signature for runnables that are triggered by time (TimingEvent) or data reception.

Caution: For intra-ECU communication, the DataReceivedEvent is fired after each completed write operation to the shared data. In case of implicit access, write operation is considered to be completed when the runnable ends. While for inter-ECU communication, the DataReceivedEvent is fired by the RTE after a callback from COM due to data reception. Over a physical network, 'data' is commonly transmitted periodically and hence not only will the latency and jitter of DataReceivedEvents vary depending on whether a configuration uses intra or inter-ECU communication, but also the number and frequency of these RTEEvents may change significantly. This means that a TimingEvent should be used to periodically activation of a runnable rather than relying on the periodic transmission of data.

5.7.5.7 DataSendCompletedEvent

Purpose: A runnable entity triggered by the RTE to receive and process transmit acknowledgment notifications.

Signature: **[rte_sws_1137]**
`void <name>([[IN Rte_Instance <instance>]])`

Notes: The runnable entity triggered by a DataSendCompletedEvent RTEEvent should use the Rte.Feedback API to actually receive the status of the acknowledgement.

5.7.6 Reentrancy

A runnable entity is declared within a software-component type. The RTE ensures that concurrent activation of same instance of a runnable entity is only allowed if the runnables attribute "canBeInvokedConcurrently" is set to TRUE (see Section 4.2.5).

When a software-component is multiply instantiated each separate instance has its own instance of the runnable entities in the software-component. Whilst instances of a

software-component are independent, the runnable entities instances share the same code (`rte_sws_2017`).

Example 5.21

Consider a component `c1` with runnable entity `re1` and entry point `ep` that is instantiated twice on the same ECU.

The two instances of `c1` each has a separate *instance* of `re1`. Software-component instances are scheduled independently and therefore each instance of `re1` could be concurrently executing `ep`.

The potential for concurrent execution of runnable entities when multiple instances of a software-component are created means that each entry point should be reentrant.

[rte_sws_3749] The RTE has to reject configurations where multiple instantiation of an AUTOSAR SW-Cs is requested and the associated attribute "supportsMultipleInstantiation" is not set to TRUE.

5.8 RTE Lifecycle API Reference

This section documents the API functions used to start and stop the RTE. RTE Lifecycle API functions are not invoked from AUTOSAR software-components – instead they are invoked from other basic software module(s).

5.8.1 Rte_Start

Purpose: Initialize the RTE itself.

Signature: **[rte_sws_2569]**
`Std_ReturnType Rte_Start(void)`

Existence: **[rte_sws_1309]** The `Rte_Start` API is always created.

Description: `Rte_Start` is intended to allocate and initialise system resources and communication resources used by the RTE.

[rte_sws_ext_2582] `Rte_Start` shall be called only once by the `EcuStateManager` after the basic software modules required by RTE are initialized. These modules include:

- OS
- COM
- memory services

The `Rte_Start` API shall not be invoked from AUTOSAR software components.

[rte_sws_2585] `Rte_Start` shall return within finite execution time – it must not enter an infinite loop.

`Rte_Start` may be implemented as a function or a macro.

Return Value: If the allocation of a resource fails, `Rte_Start` shall return with an error.

- **[rte_sws_1261]** `RTE_E_OK` – No error occurred.
- **[rte_sws_1262]** `RTE_E_LIMIT` – An internal limit has been exceeded. The allocation of a required resource has failed.

Notes: `Rte_Start` is declared in the lifecycle header file `Rte_Main.h`. The initialization of AUTOSAR software-components takes place after the termination of `Rte_Start` and is triggered by a mode change event on entering run state.

5.8.2 Rte_Stop

Purpose: finalize the RTE itself

Signature: **[rte_sws_2570]**
`Std_ReturnType Rte_Stop(void)`

Existence: **[rte_sws_1310]** The `Rte_Stop` API is always created.

Description: `Rte_Stop` is used to finalize the RTE itself. This service releases all system and communication resources allocated by the RTE.

[rte_sws_ext_2583] `Rte_Stop` shall be called by the `EcuStateManager` before the basic software modules required by RTE are shut down. These modules include:

- OS
- COM
- memory services

`Rte_Stop` shall not be called by an AUTOSAR software component.

[rte_sws_2584] `Rte_Stop` shall return within finite execution time.

`Rte_Stop` may be implemented as a function or a macro.

Return Value:

- **[rte_sws_1259]** `RTE_E_OK` – No error occurred.
- **[rte_sws_1260]** `RTE_E_LIMIT` – a resource could not be released.

Notes: `Rte_Stop` is declared in the lifecycle header file `Rte_Main.h`.

5.9 RTE Call-backs Reference

This section documents the call-backs that are generated by the RTE that must be invoked by other components, such as the communication service, and therefore must have a well-defined name and semantics.

[rte_sws_1165] A call-back implementation created by the RTE generator is not permitted to block.

Requirement `rte_sws_1165` serves to constrain RTE implementations so that all implementations can work with all basic software.

5.9.1 RTE-COM Message Naming Conventions

The COM signals used for communication are defined in the meta-model (Section B).

[rte_sws_3007] The RTE shall initiate an inter-ECU transmission using the COM API with the handle id of the corresponding COM signal for primitive data element `rte_sws_in_0063`.

[rte_sws_3008] The RTE shall initiate an inter-ECU transmission using the COM API with the handle id of the corresponding COM signal group for complex data element or operation arguments `rte_sws_in_0064`.

5.9.2 Communication Service Call-backs

Purpose: Implement the call-back functions that AutoSAR COM invokes as a result of inter-ECU communication, where:

- A data item/event is ready for reception by a receiver.
- A transmission acknowledgment shall be routed to a sender.
- An operation shall be invoked by a server.
- The result of an operation is ready for reading by a client.

Signature: **[rte_sws_3000]**

```
void <CallbackRoutineName> (void);
```

Where `<CallbackRoutineName>` is the name of the call-back function (refer to Section 5.9.3 for details on the naming convention).

Description: Prototypes for the call-back `<CallbackRoutineName>` provided by AutoSAR COM.

Return Value: No return value : `void`

5.9.3 Naming convention of CallbackRoutineName

In the following table, the naming convention of `<CallBackRoutineName>` are defined:

Calling Situation	callbackRoutineName	Comments
A primitive data item/event is ready for reception by a receiver.	[rte_sws_3001] Rte_COMCbk_<sn>	<sn> is the name of the COM signal. This callback function indicates that the signal of the primitive data item/event or the single argument of an operation is ready for reception.
A transmission acknowledgment of a primitive data item/event shall be routed to a sender.	[rte_sws_3002] Rte_COMCbkTack_<sn>	“Tack” is literal text indicating transmission acknowledgment. This callback function indicates that the signal of the primitive data item/event is already handed over by COM to the PDU router.
A transmission error notificatoin of a primitive data item/event shall be routed to a sender.	[rte_sws_3775] Rte_COMCbkTErr_<sn>	“TErr” is literal text indicating transmission error. This callback function indicates that an error occurred when the signal of the primitive data item/event was handed over by COM to the PDU router.
A signal invalidation of a primitive data item shall be routed to a receiver.	[rte_sws_2612] Rte_COMCbkInv_<sn>	“Inv” is literal text indicating signal invalidation. This callback function indicates that COM has received a signal and parsed it as “invalid”.
A signal of a primitive data item is outdated. No new data is available.	[rte_sws_2610] Rte_COMCbkTOut_<sn>	“TOut” is literal text indicating signal time out. This callback function indicates that the aliveTime-out after the last successful reception of the signal of the primitive data item/event has expired (data element outdated).
A complex data item/event or the arguments of an operation is ready for reception by a receiver.	[rte_sws_3004] Rte_COMCbk_<sg>	<sg> is the name of the COM signal group, which contains all the signals of the complex data item/event or an operation. This callback function indicates that the signals of the complex data item/event or the arguments of an operation are ready for reception.
A transmission acknowledgment of a complex data item/event shall be routed to a sender.	[rte_sws_3005] Rte_COMCbkTack_<sg>	“Tack” is literal text indicating transmission acknowledgment. This callback function indicates that the signals of the complex data item/event is already handed over by COM to the PDU router.

Calling Situation	callbackRoutineName	Comments
A transmission error notificatoion of a complex data item/event shall be routed to a sender.	[rte_sws_3776] Rte_COMCbkJTErr_<sg>	“TErr” is literal text indicating transmission error. This callback function indicates that an error occurred when the signal of the complex data item/event was handed over by COM to the PDU router.
A signal group invalidation of a composite data item shall be routed to a receiver.	[rte_sws_5065] Rte_COMCbkJInv_<sg>	“Inv” is literal text indicating signal group invalidation. This callback function indicates that COM has received a signal group and parsed it as “invalid”.
A signal group of a complex data item is outdated. No new data is available.	[rte_sws_2611] Rte_COMCbkJTOut_<sg>	“TOut” is literal text indicating signal time out. This callback function indicates that the aliveTimeout after the last successful reception of the signal group carrying the complex data item has expired (data element outdated).

Table 5.7: RTE COM Callback Function Naming Conventions

Where:

- <sn> is a COM signal name.
- <sg> is a COM signal group name.

5.10 VFB Tracing Reference

The RTE’s “VFB Tracing” functionality permits the monitoring of AUTOSAR signals as they are sent and received across the VFB.

The RTE operates in at least two builds (some implementations may provide more than two builds). One does not enable VFB tracing whereas the other can be configured to trace some or all “interesting events”.

[rte_sws_1327] The RTE generator shall support a build where no VFB events are traced.

[rte_sws_1328] The RTE generator shall support a build that traces (configured) VFB events.

The RTE generator’s ‘trace’ build is enabled or disabled through definitions in the RTE Configuration file `rte_sws_1322` and `rte_sws_1323`. Note that this ‘trace’ build is intended to enable tracing of software components and not the RTE itself.

5.10.1 Principle of Operation

The “VFB Tracing” mechanism is designed to offer a lightweight means to monitor the interactions of AUTOSAR software-components with the VFB.

The VFB tracing is implemented by a series of “hook” functions that are invoked automatically by the generated RTE when “interesting events” occur. Each hook function corresponds to a single event.

The supported trace events are defined in Section 5.10.2. A mechanism is described in Section 5.10.3 for configuring which of the many potential trace events are of interest.

5.10.2 Trace Events

5.10.2.1 RTE API Trace Events

RTE API trace events occur when an AUTOSAR software-component interacts with the generated RTE API. For implicit S/R communication, however, tracing is not supported.

5.10.2.1.1 RTE API Start

Description: RTE API Start is invoked by the RTE when an API call is made by a component.

Signature: [rte_sws_1238]

```
void Rte_<api>Hook_<c>_<ap>_Start
    ([const Rte_CDS_<c>>*, ]<param>)
```

Where <api> is the RTE API Name (Write, Call, etc.), <c> is the component type name and <ap> the access point name (e.g. port and data element or operation name, exclusive area name, etc.). The parameters of the API are the same as the corresponding RTE API. As with the API itself, the instance handle is included if and only if the software component’s SupportsMultipleInstantiation (rte_sws_in_0004) attribute is set to true. Note that `Rte_Instance` cannot be used directly, as there will be pointers to multiple components’ structure types within the single VFB Tracing header file, and `Rte_Instance` would therefore be ambiguous.

5.10.2.1.2 RTE API Return

Description: RTE API Return is a trace event that is invoked by the RTE just before an API call returns control to a component.

Signature: [rte_sws_1239]

```
void Rte_<api>Hook_<c>_<ap>_Return
    ([const Rte_CDS_<c>*, ]<param>)
```

Where `<api>` is the RTE API Name (Write, Call, etc.), `<c>` is the component type name and `<ap>` the access point name (e.g. port and data element or operation name, exclusive area name, etc.). The parameters of the API are the same as the corresponding RTE API and contain the values of OUT and INOUT parameters on exit from the function.

As with the API itself, the instance handle is included if and only if the software component's `SupportsMultipleInstantiation` (`rte_sws.in_0004`) attribute is set to true. Note that `Rte_Instance` cannot be used directly, as there will be pointers to multiple components' structure types within the single VFB Tracing header file, and `Rte_Instance` would therefore be ambiguous.

5.10.2.2 COM Trace Events

COM trace events occur when the generated RTE interacts with the AUTOSAR communication service.

5.10.2.2.1 Signal Transmission

Description: A trace event indicating a transmission request of an Inter-ECU signal or signal group by the RTE. Invoked by the RTE just before `Com_SendSignal` or `Com_UpdateShadowSignal` is invoked.

Signature: **[rte_sws_1240]**

```
void Rte_ComHook_<signalName>_SigTx(<data>)
```

Where `<signalName>` is the COM signal name and `<data>` a pointer to the signal data to be transmitted.

5.10.2.2.2 Signal Reception

Description: A trace event indicating a successful attempt to read an Inter-ECU signal by the RTE. Invoked by the RTE after return from `Com_ReceiveSignal` or `Com_ReceiveShadowSignal`.

Signature: **[rte_sws_1241]**

```
void Rte_ComHook_<signalName>_SigRx(<data>)
```

Where `<signalName>` is the COM signal name and `<data>` a pointer to the signal data received.

5.10.2.2.3 Signal Invalidation

Description: A trace event indicating a signal invalidation request of an Inter-ECU signal or signal group by the RTE. Invoked by the RTE just before `Com_InvalidateSignal` or `Com_InvalidateShadowSignal` is invoked.

Signature: **[rte_sws_3814]**
`void Rte_ComHook_<signalName>_SigIv(void)`

Where `<signalName>` is the COM signal name.

5.10.2.2.4 COM Callback

Description: A trace event indicating the start of a COM call-back. Invoked by generated RTE code on entry to the COM call-back.

Signature: **[rte_sws_1242]**
`void Rte_ComHook<Event>_<signalName>(void)`

Where `<signalName>` is the name of the COM signal or signal group and `<Event>` indicates the callback type and can take the values “Inv” for an invalidation callback, “TOut” for a timeout callback, “Tack” for a transmission acknowledgement callback, or “TErr” for a transmission error callback.

5.10.2.3 OS Trace Events

OS trace events occur when the generated RTE interacts with the AUTOSAR operating system.

5.10.2.3.1 Task Activate

Description: A trace event that is invoked by the RTE immediately prior to the activation of a task containing runnable entities.

Signature: **[rte_sws_1243]**
`void Rte_Task_Activate(TaskType task)`

Where `task` is the OS’s handle for the task.

5.10.2.3.2 Task Dispatch

Description: A trace event that is invoked immediately an RTE generated task (containing runnable entities) has commenced execution.

Signature: **[rte_sws_1244]**
`void Rte_Task_Dispatch(TaskType task)`

Where `task` is the OS's handle for the task.

5.10.2.3.3 Set OS Event

Description: A trace event invoked immediately before generated RTE code attempts to set an OS Event.

Signature: **[rte_sws_1245]**
`void Rte_Task_SetEvent(TaskType task, EventMaskType ev)`

Where `task` is the OS's handle for the task for which the event is being set and `ev` the OS event mask.

5.10.2.3.4 Wait OS Event

Description: Invoked immediately before generated RTE code attempts to wait on an OS Event. This trace event does *not* indicate that the caller has suspended execution since the OS call may immediately return if the event was already set.

Signature: **[rte_sws_1246]**
`void Rte_Task_WaitEvent(TaskType task, EventMaskType ev)`

Where `task` is the OS's handle for the task (that is waiting for the event) and `ev` the OS event mask.

5.10.2.3.5 Received OS Event

Description: Invoked immediately after generated RTE code returns from waiting on an event.

Signature: **[rte_sws_1247]**
`void Rte_Task_WaitEventRet(TaskType task,
EventMaskType ev)`

Where `task` is the OS's handle for the task (that was waiting for an event) and `ev` the event mask indicating the received event.

Note that not all of the trace events listed above may be available for a given input configuration. For example if a task is activated by a schedule table, it is activated by the OS rather than by the RTE, hence no trace hook function for task activation can be invoked by the RTE.

5.10.2.4 Runnable Entity Trace Events

Runnable entity trace events occur when a runnable entity is started.

5.10.2.4.1 Runnable Entity Invocation

Description: Event invoked by the RTE just before execution of runnable entry starts via its entry point. This trace event occurs after any copies of data elements are made to support the `Rte_IRead` API Call.

Signature: **[rte_sws_1248]**

```
void Rte_Runnable_<c>_<reName>_Start
    ([const RTE_CDS_<c>*])
```

Where `<c>` is the SW-C type name and `reName` the runnable entity name. The instance handle is included if and only if the software component's `SupportsMultipleInstantiation` (`rte_sws.in_0004`) attribute is set to true. Note that `Rte_Instance` cannot be used directly, as there will be pointers to multiple components' structure types within the single VFB Tracing header file, and `Rte_Instance` would therefore be ambiguous.

5.10.2.4.2 Runnable Entity Termination

purpose: Event invoked by the RTE immediately execution returns to RTE code from a runnable entity. This trace event occurs before any write-back of data elements are made to support the `Rte_IWrite` API Call.

Signature: **[rte_sws_1249]**

```
void Rte_Runnable_<c>_<reName>_Return
    ([const Rte_CDS_<c>*])
```

Where `<c>` is the SW-C type name and `reName` the runnable entity name. The instance handle is included if and only if the software component's `SupportsMultipleInstantiation` (`rte_sws.in_0004`) attribute is set to true. Note that `Rte_Instance` cannot be used directly, as there will be pointers to multiple components' structure types within the single VFB Tracing header file, and `Rte_Instance` would therefore be ambiguous.

5.10.3 Configuration

The VFB tracing mechanism works by the RTE invoking the tracepoint *hook* function whenever the tracing event occurs.

The support trace events and their hook function name and signature are defined in Section 5.10.2. There are many potential trace events and it is likely that only a few will be of interest at any one time. Therefore The RTE generator supports a mechanism to configure which trace events are of interest.

In order to minimise RTE Overheads, trace events that are not enabled should have no run-time effect on the generated system. This is achieved through generated code within the VFB Tracing Header File (see Section 5.3.5) and the user supplied definitions from the RTE Configuration Header file (see Section 5.3.6).

The definition of trace event hook functions is contained within user code. If a definition is encapsulated within a `#if` block, as follows, the definition will automatically be omitted when the trace event is disabled.

```
1  #if !defined(<trace event>)
2  void <trace event>(<params>)
3  {
4      /* Function definition */
5  }
6  #endif
```

The configuration of which individual trace events are enabled is entirely under the control of the user via the definitions included in the RTE Configuration header file.

5.10.4 Interaction with Object-code Software-Components

VFB tracing is only available during the “RTE Generation” phase `rte_sws_1319` and therefore hook functions never appear in an application header file created during “RTE Contract” phase. However, object-code software-components are compiled against the “RTE Contract” phase header and can therefore only trace events that are inserted into the generated RTE. In particular they cannot trace events that require invocation of hook functions to be inserted into the API mapping such as the `Rte_Pim` API. However, many trace events are applicable to object-code software-components including trace events related to the explicit communication API, to task activity and for runnable entity start and stop.

This approach means that the external interactions of the object-code software-component can be monitored without requiring modification of the delivered object-code and without revealing the internal activity of the software-component. The approach is therefore considered to be consistent with the desire for IP protection that prompts delivery of a software-component as object-code.

6 RTE ECU Configuration

The RTE provides the glue layer between the application SW-Components and the Basic Software thus enabling several application SW-Components to be integrated on one ECU. The RTE layer is shown in figure 6.1.

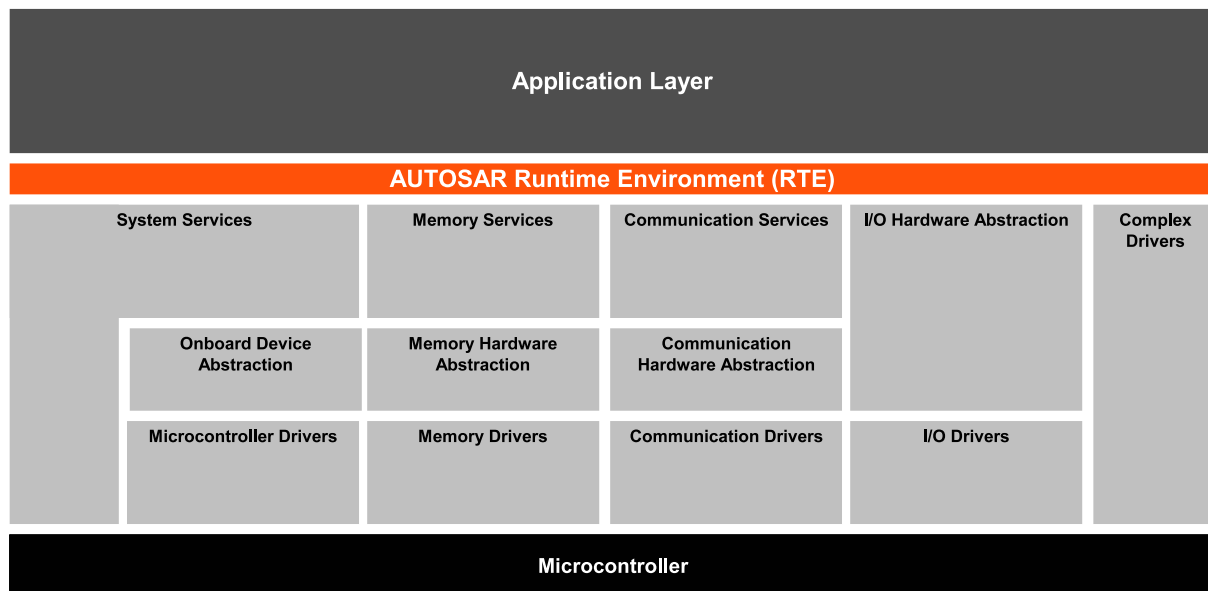


Figure 6.1: ECU Architecture RTE

The RTE is included twice in the development methodology of SW-Components. In the RTE Contract phase only the SW-Component description is used as an input. The configuration parameters defined in this section are used for the RTE Generation phase.

The overall structure of the RTE configuration parameters is shown in figure 6.2. It has to be distinguished between the configuration parameters for the RTE generator and the configuration parameters for the generated RTE itself.

Most of the information needed to generate an RTE is already available in the ECU extract of the System Description. From this extract also the links to the SW-Component descriptions and ECU Resource description are available. So only additional information not covered by the three aforementioned formats needs to be provided by the ECU Configuration description.

To additionally allow the most flexibility and freedom in the implementations of the RTE, only configuration parameters which are common to all implementations are standardized in the ECU Configuration Parameter definition. Any additional configuration parameters which might be needed to configure a full functional RTE have to be specified using the vendor specific parameter definition mechanism described in the ECU Configuration specification document [7].

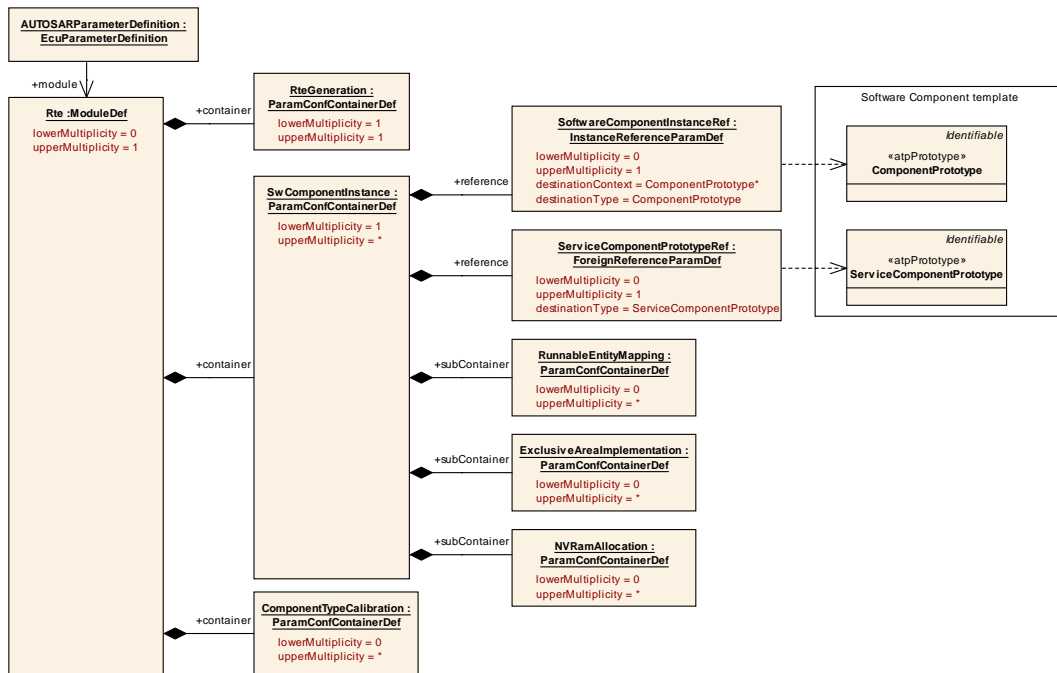


Figure 6.2: RTE configuration overview

The configuration of the RTE is structured in the following groups:

- RteGeneration in section 6.1
- ImplementationSelection in section 6.2.1
- RunnableEntityMapping in section 6.2.2
- ExclusiveAreaImplementation in section 6.2.3
- NVRamAllocation in section 6.2.4
- CalprmComponentInstance in section 6.3

Module Name	Rte	
Module Description	Configuration of the Rte (Runtime Environment) module.	
Included Containers		
Container Name	Multiplicity	Scope / Dependency
ComponentType Calibration	0..*	Specifies for each CalprmComponentType or AtomicSoftwareComponentType whether calibration is enabled.
RteGeneration	1	This container holds the parameters for the configuration of the RTE Generation.

Container Name	Multiplicity	Scope / Dependency
SwComponentInstance	1..*	<p>Representation of one SW-Component instance located on the to be configured ECU. All subcontainer configuration aspects are in relation to this SW-Component instance.</p> <p>The SwComponentInstance can be either an ApplicationSoftwareComponentInstance or a ServiceComponentInstance.</p>

6.1 RTE Generation Parameters

The parameters in the container `RteGeneration` are used to configure the RTE generator. They all need to be defined during pre-compile time.

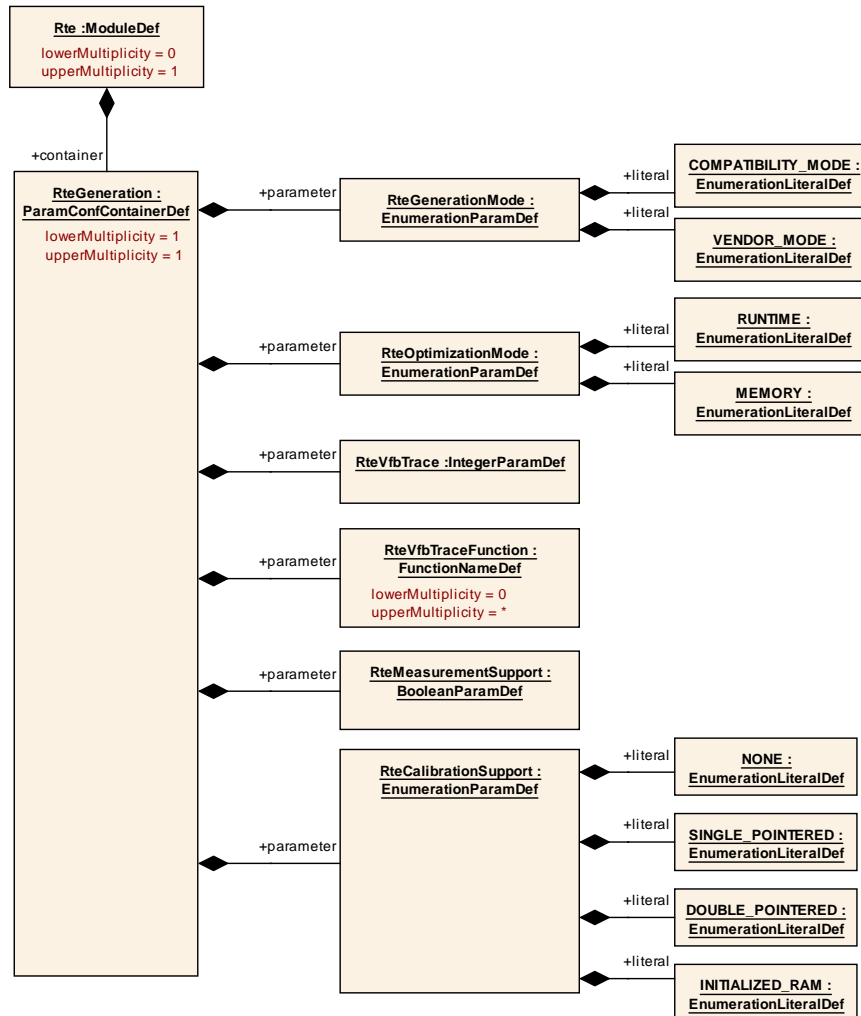


Figure 6.3: RTE generation parameters

RteGeneration

SWS Item	
Container Name	RteGeneration
Description	This container holds the parameters for the configuration of the RTE Generation.
Configuration Parameters	

Name	RteCalibrationSupport		
Description	The RTE generator shall have the option to switch off support for calibration for generated RTE code. This option shall influence complete RTE code at once.		
Multiplicity	1		
Type	EnumerationParamDef		
Range	DOUBLE_POINTERED		
	INITIALIZED_RAM		
	NONE		
	SINGLE_POINTERED		
Configuration Class	Pre-compile time	X	All Variants
	Link time	–	
	Post-build time	–	
Scope / Dependency			

Name	RteGenerationMode		
Description	Switch between the two available generation modes of the RTE generator.		
Multiplicity	1		
Type	EnumerationParamDef		
Range	COMPATIBILITY_MODE		
	VENDOR_MODE		
Configuration Class	Pre-compile time	X	All Variants
	Link time	–	
	Post-build time	–	
Scope / Dependency			

Name	RteMeasurementSupport		
Description	The RTE generator shall have the option to switch off support for measurement for generated RTE code. This option shall influence complete RTE code at once.		
Multiplicity	1		
Type	BooleanParamDef		
Default Value			
Configuration Class	Pre-compile time	X	All Variants
	Link time	–	
	Post-build time	–	
Scope / Dependency			

Name	RteOptimizationMode		
Description	Switch between the two available optimization modes of the RTE generator.		
Multiplicity	1		
Type	EnumerationParamDef		
Range	MEMORY		
	RUNTIME		
Configuration Class	Pre-compile time	X	All Variants
	Link time	–	
	Post-build time	–	
Scope / Dependency			

Name	RteVfbTrace {RTE_VFB_TRACE}		
Description	The RTE generator shall globally enable VFB tracing when RTE_VFB_TRACE is defined in the RTE configuration header file as a non-zero integer. The RTE generator shall globally disable VFB tracing when RTE_VFB_TRACE is defined in the RTE configuration header file as 0.		
Multiplicity	1		
Type	IntegerParamDef		
Default Value			
Configuration Class	Pre-compile time	X	All Variants
	Link time	–	
	Post-build time	–	
Scope / Dependency			
Name	RteVfbTraceFunction		
Description	The RTE generator shall enable VFB tracing for a given hook function when there is a #define in the RTE configuration header file for the hook function name and tracing is globally enabled. Example: #define Rte_WriteHook_i1_p1_a_Start		
Multiplicity	0..*		
Type	FunctionNameDef		
Default Value			
Configuration Class	Pre-compile time	X	All Variants
	Link time	–	
	Post-build time	–	
Scope / Dependency			
No Included Containers			

6.2 Handling of Software Component instances

When entities of Software Components are to be configured there is the need to actually address the instances of the `AtomicSoftwareComponentType`. AUTOSAR defines the *instance reference* as the mechanism to allow this addressing.

The special semantics of the `InstanceReferenceDef`¹ does allow to address each "instance" of that `ComponentPrototype` in the Software Component template. Since the whole vehicle is described as one `top-level composition` using the Software Component template [17], each actual instance of any `AtomicSoftwareComponentType` from the VFB view can be addressed using this mechanism.

Since the Service Components are not part of the VFB they can not be referenced using the `SoftwareComponentInstanceRef`. Service Components are referenced using the foreign reference `ServiceComponentPrototypeRef`.

In figure 6.2 the container `SwComponentInstance` is shown which contains the instance reference `SoftwareComponentInstanceRef` and the foreign reference `ServiceComponentPrototypeRef`. These references are used to unambiguously identify each instance of a `AtomicSoftwareComponentType` within the ECU.

SwComponentInstance

SWS Item			
Container Name	SwComponentInstance		
Description	Representation of one SW-Component instance located on the to be configured ECU. All subcontainer configuration aspects are in relation to this SW-Component instance. The SwComponentInstance can be either an ApplicationSoftwareComponentInstance or a ServiceComponentInstance.		
Configuration Parameters			
Name	ImplementationRef		
Description	The Implementation which is assigned to the ComponentPrototype.		
Multiplicity	0..1		
Type	Foreign reference to SwImplementation		
Configuration Class	Pre-compile time	X	All Variants
	Link time	–	
	Post-build time	–	
Scope / Dependency			

¹The `InstanceReferenceDef` is based on the `<<instanceRef>>` mechanism introduced in the "Template UML Profile and Modeling Guide" document [16].

Name	ServiceComponentPrototypeRef		
Description	Reference to the ServiceComponentPrototype representing an particular service on this ECU. If ServiceComponentPrototypeRef is specified there shall not be the SoftwareComponentInstanceRef specified.		
Multiplicity	0..1		
Type	Foreign reference to ServiceComponentPrototype		
Configuration Class	Pre-compile time	–	
	Link time	–	
	Post-build time	–	
Scope / Dependency			

Name	SoftwareComponentInstanceRef		
Description	Reference to a SW-Component ComponentPrototype. If SoftwareComponentInstanceRef is specified there shall not be the ServiceComponentPrototypeRef specified. Semantic Constraint: Only ComponentPrototypes which have an AtomicSoftwareComponentType as <<isOfType>> shall be referenced here.		
Multiplicity	0..1		
Type	Instance reference to ComponentPrototype context: Component Prototype*		
Configuration Class	Pre-compile time	X	All Variants
	Link time	–	
	Post-build time	–	
Scope / Dependency			
Included Containers			
Container Name	Multiplicity	Scope / Depedency	
ExclusiveArea Implementation	0..*		
NVRamAllocation	0..*		
RunnableEntityMapping	0..*		

The container `SwComponentInstance` collects all the configuration information related to one specific instance of a `AtomicSoftwareComponentType`. The individual aspects will be described in the next sections.

6.2.1 Selection of SW-Component Implementation

During the system development there is no need to select the actual implementation which will be later integrated on one ECU. Therefore the ECU Extract of System Description may not contain the selection information yet.

In the SW-Component template an `SwcImplementation` is always associated – via an `InternalBehavior` – with an `AtomicSoftwareComponentType`.

In theory it is possible to have different `SwcImplementation` (and implicitly different `InternalBehavior`) provided for each instance of an `AtomicSoftwareComponentType`. This is currently not supported by the RTE, so the following restriction applies:

For each `ComponentPrototype` of the same `AtomicSoftwareComponentType` the identical `SwcImplementation` and `InternalBehavior` shall be configured.

The mapping of `SwcImplementation` to `ComponentPrototype` done using the two references `SoftwareComponentInstanceRef` and `ImplementationRef` (see figure 6.4).

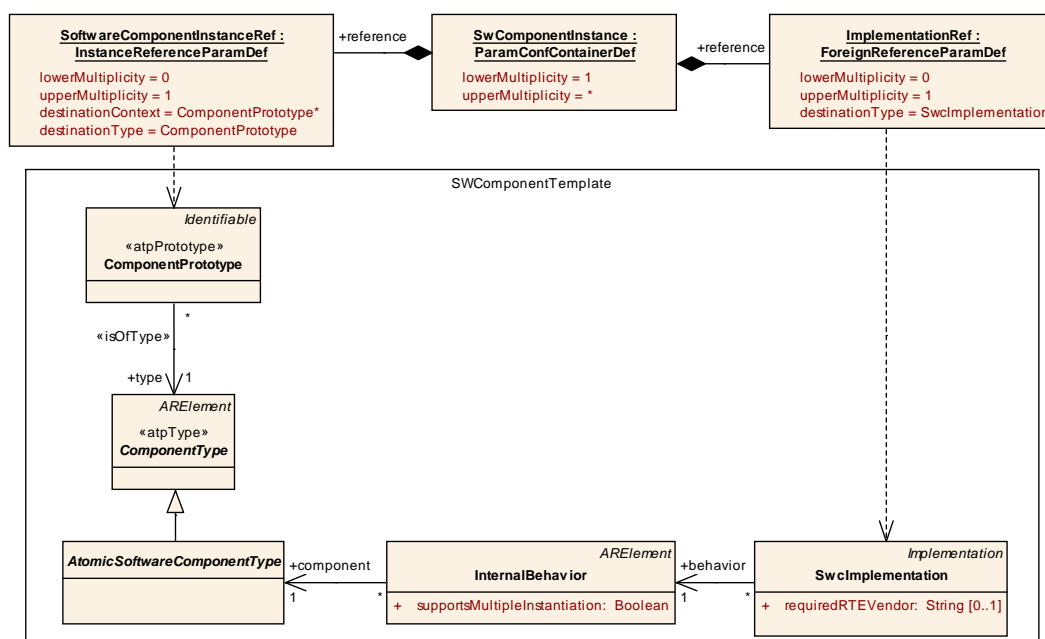


Figure 6.4: Selection of the Implementation for an AtomicSoftwareComponentType

6.2.2 Runnable Entity to task mapping

One of the major fragments of the RTE configuration is the mapping of SW-Component's `RunnableEntity`s to OS Tasks. The parameters defined to achieve this are shown in figure 6.5.

The mapping is based on the `RTEEvent` because it is the source of the activation. For each `RunnableEntity` which belongs to a SW-Component instance mapped on the ECU there needs to be a mapping container specifying how this `RunnableEntity` should be handled. If the `RunnableEntity` is a server-runnable and shall be executed in the context of the caller (i.e. using a direct function call) the element `RunnableEntityMapping` still shall be provided to indicate that this `RTEEvent` has been considered in the mapping, but no further parameters or references are required (e.g. `MappedToTaskRef` can be left out).




Figure 6.5: RTE runnable entity to task mapping

One major constraint is posed by the `canBeInvokedConcurrently` attribute of each `RunnableEntity` because data consistency issues have to be considered.

The `MappedToTaskRef` `OsTask` is part of the ECU Configuration description, so a plain `ReferenceDef` can be used to establish the link.

Another important parameter is the `PositionInTask` which provides an order of `RunnableEntity`s within the associated `OsTask`. When the task is executed periodically the `PositionInTask` parameter defines the order of execution within the test. When the task is used to define a context for event activated `RunnableEntity`s the `PositionInTask` parameter defines the order of evaluation which actual `RunnableEntity` shall be executed. Thus providing means to define a deterministic delay between the beginning of execution of the task and the actual execution of the `RunnableEntity`'s code.

When an `OsEvent` is used to activate the `OsTask` the reference `UsedOsEventRef` specifies which `OsEvent` is used.

RunnableEntityMapping

SWS Item	
Container Name	RunnableEntityMapping
Description	Maps a RunnableEntity onto one OS Task based on the activating RTEEvent. Even if a RunnableEntity is executed in the caller's context this RunnableEntityMapping shall be specified, but no MappedToTask and PositionInTask elements given.
Configuration Parameters	

Name	ActivationOffset		
Description	Activation offset in seconds.		
Multiplicity	0..1		
Type	FloatParamDef		
Default Value			
Configuration Class	Pre-compile time	X	All Variants
	Link time	–	
	Post-build time	–	
Scope / Dependency			

Name	MappedToTaskRef		
Description	Reference to the OsTask the RunnableEntity is mapped to. If no reference to the OsTask is specified the RunnableEntity is executed in the context of the caller.		
Multiplicity	0..1		
Type	Reference to OsTask		
Configuration Class	Pre-compile time	X	All Variants
	Link time	–	
	Post-build time	–	
Scope / Dependency			

Name	PositionInTask		
Description	Each RunnableEntities mapped to an OsTask has a specific position within the task execution. For periodic activation this is the order of execution. For event driver activation this is the order of evaluation which actual RunnableEntity has to be executed.		
Multiplicity	0..1		
Type	IntegerParamDef		
Default Value			
Configuration Class	Pre-compile time	X	All Variants
	Link time	–	
	Post-build time	–	
Scope / Dependency			

Name	RTEEventRef		
Description	Reference to the description of the RTEEvent which is pointing to the RunnableEntity being mapped. This allows a fine grained mapping of RunnableEntites based on the activating RTEEvent.		
Multiplicity	1		
Type	Foreign reference to RTEEvent		
Configuration Class	Pre-compile time	X	All Variants
	Link time	–	
	Post-build time	–	
Scope / Dependency			
Name	UsedOsEventRef		
Description	If an OsEvent is used to activate the RunnableEntity it shall be referenced here.		
Multiplicity	0..1		
Type	Reference to OsEvent		
Configuration Class	Pre-compile time	X	All Variants
	Link time	–	
	Post-build time	–	
Scope / Dependency			
No Included Containers			

There are some constraints which do apply when actually mapping the RunnableEntity to an OsTask:

[rte_sws_5082] The following restrictions apply to RTEEvents which are used to activate RunnableEntity. OsEvents that are used to wakeUpFromWaitPoint shall not be included in the mapping.

When a wakeUpFromWaitPoint is occurring the RunnableEntity resumes its execution in the context of the originally activated OsTask.

[rte_sws_5083] If the canBeInvokedConcurrently (rte_sws_in_0072) flag of the RunnableEntity is false all mappings of that RunnableEntity have to point to the same OsTask.

6.2.3 Exclusive Area implementation

The RTE Generator can be configured to implement a different data consistency mechanism for each ExclusiveArea defined for a SW-Component.

In figure 6.6 the configuration of the actually selected data consistency mechanism is shown.

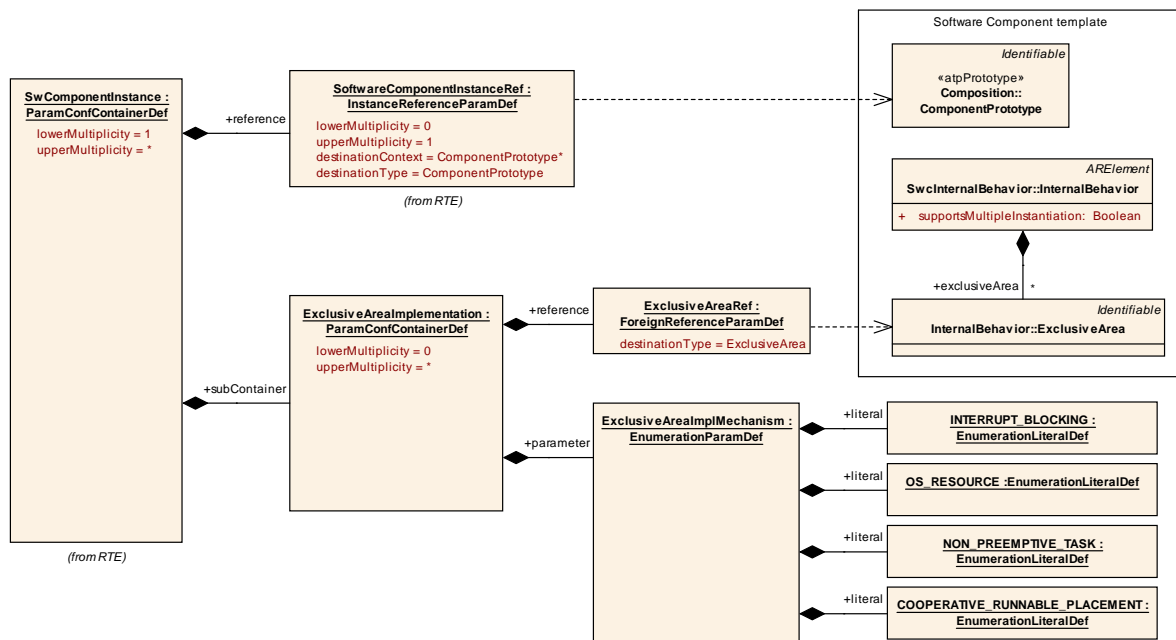


Figure 6.6: Configuration of the ExclusiveArea implementation

ExclusiveAreaImplementation

SWS Item			
Container Name	ExclusiveAreaImplementation		
Description	Specifies the implementation to be used for the data consistency of this ExclusiveArea.		
Configuration Parameters			
Name	ExclusiveAreaImplMechanism		
Description	To be used implementation mechanism for the specified ExclusiveArea.		
Multiplicity	1		
Type	EnumerationParamDef		
Range	COOPERATIVE_RUNNABLE_PLACEMENT		
	INTERRUPT_BLOCKING		
	NON_PREEMPTIVE_TASK		
	OS_RESOURCE		
Configuration Class	Pre-compile time	X	All Variants
	Link time	–	
	Post-build time	–	
Scope / Dependency			

Name	ExclusiveAreaRef		
Description	Reference to the ExclusiveArea.		
Multiplicity	1		
Type	Foreign reference to ExclusiveArea		
Configuration Class	Pre-compile time	X	All Variants
	Link time	–	
	Post-build time	–	
Scope / Dependency			
No Included Containers			

6.2.4 NVRam Allocation

The configuration of the NVRam access does involve several templates, because it closes the gap between the SW-Components, the Services and the BSW Modules.

In figure 6.7 the related information from the SW-Component template is shown.

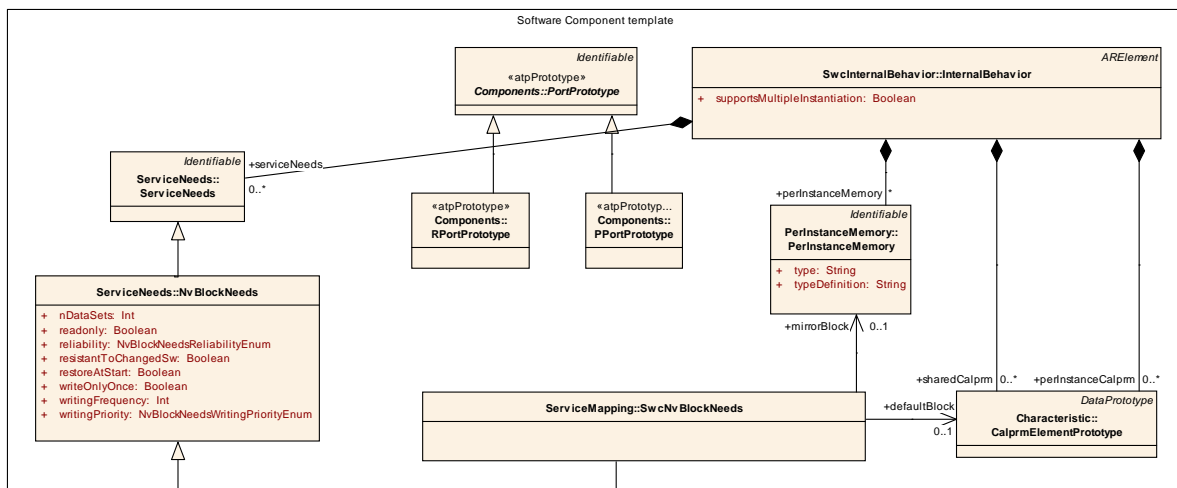


Figure 6.7: SW-Component information of NVRam Service needs

In figure 6.8 the ECU Configuration part of the NVRam allocation is shown. It relates the SW-Components `NVRAMMapping` information with the NVRam Managers `NvmBlockDescriptor` and the linker symbols of the RAM and ROM sections to be used.

NVRamAllocation

SWS Item	
Container Name	NVRamAllocation
Description	Specifies the relationship between the SW-Components <code>NVRAMMapping</code> / NVRAM needs and the NvM module configuration.
Configuration Parameters	

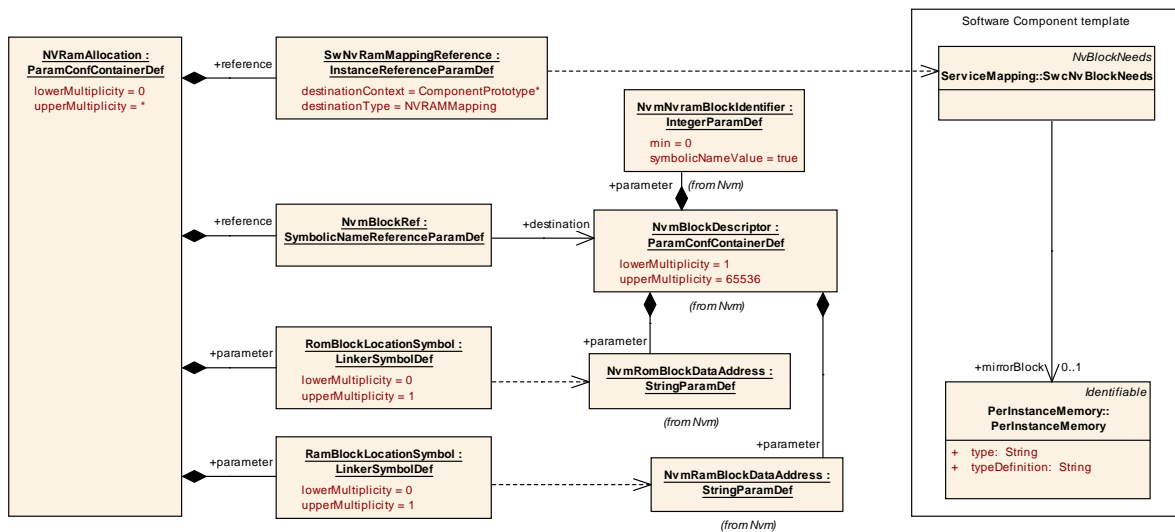


Figure 6.8: ECU Configuration of the NVRam Service

Name	NvmBlockRef		
Description	Reference to the used NvmBlock for storage of the NVRAMMapping information.		
Multiplicity	1		
Type	Symbolic name reference to NvmBlockDescriptor		
Configuration Class	Pre-compile time	X	All Variants
	Link time	–	
	Post-build time	–	
Scope / Dependency			

Name	RamBlockLocationSymbol		
Description	This is the name of the linker object name where the NVRam Block will be mirrored by the Nvm. This symbol will be resolved into the parameter "NvmRamBlockDataAddress" from the "NvmBlockDescriptor".		
Multiplicity	0..1		
Type	LinkerSymbolDef		
Default Value			
Configuration Class	Pre-compile time	X	All Variants
	Link time	–	
	Post-build time	–	
Scope / Dependency			

Name	RomBlockLocationSymbol		
Description	This is the name of the linker object name where the NVRom Block will be accessed by the Nvm. This symbol will be resolved into the parameter "NvmRomBlockDataAddress" from the "NvmBlockDescriptor".		
Multiplicity	0..1		
Type	LinkerSymbolDef		
Default Value			
Configuration Class	Pre-compile time	X	All Variants
	Link time	-	
	Post-build time	-	
Scope / Dependency			

Name	SwNvRamMappingReference		
Description	Reference to the NVRAMMapping instance to be configured.		
Multiplicity	1		
Type	Instance reference to NVRAMMapping context: ComponentPrototype*		
Configuration Class	Pre-compile time	X	All Variants
	Link time	-	
	Post-build time	-	
Scope / Dependency			
No Included Containers			

6.3 Component Type Calibration

In the SW-Component template two places may provide calibration data: the `CalprmComponentType` and the `AtomicSoftwareComponentType` (or more precisely the subclasses of `AtomicSoftwareComponentType`). Whether the calibration is enabled for a specific `ComponentType` can be configured as shown in figure 6.9.

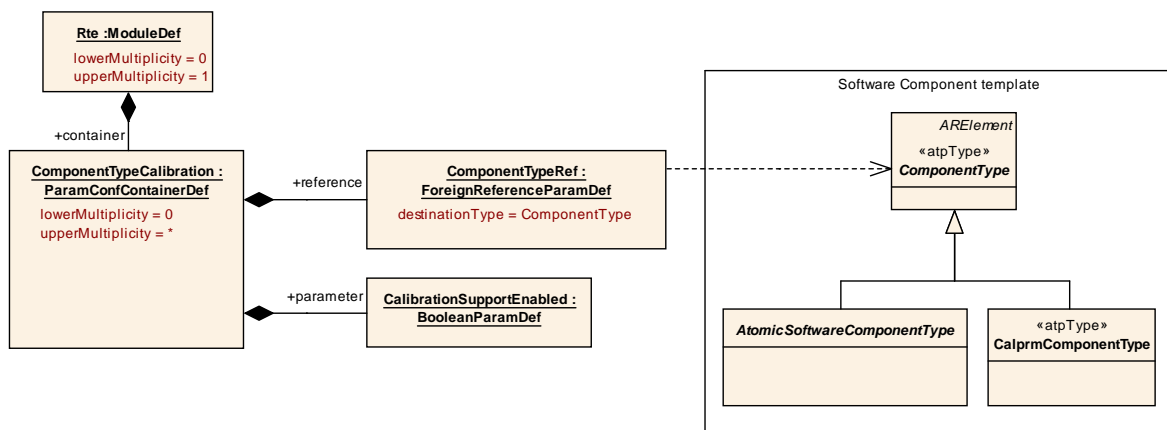


Figure 6.9: Configuration of the calibration for the `CalprmComponentType`

The foreign reference `ComponentTypeRef` identifies the `ComponentType` (which is limited to `CalprmComponentType` and `AtomicSoftwareComponentType`). The

boolean parameter `CalibrationSupportEnabled` specifies whether calibration shall be enabled for the specified `ComponentType`.

ComponentTypeCalibration

SWS Item			
Container Name	ComponentTypeCalibration		
Description	Specifies for each CalprmComponentType or AtomicSoftwareComponentType whether calibration is enabled.		
Configuration Parameters			
Name	CalibrationSupportEnabled		
Description	Enables calibration support for the specified CalprmComponentType or AtomicSoftwareComponentType.		
Multiplicity	1		
Type	BooleanParamDef		
Default Value			
Configuration Class	Pre-compile time	X	All Variants
	Link time	–	
	Post-build time	–	
Scope / Dependency			
Name	ComponentTypeRef		
Description	Reference to the CalprmComponentType or AtomicSoftwareComponentType.		
Multiplicity	1		
Type	Foreign reference to ComponentType		
Configuration Class	Pre-compile time	X	All Variants
	Link time	–	
	Post-build time	–	
Scope / Dependency			
No Included Containers			

6.4 Communication infrastructure

The configuration of the communication infrastructure (interaction of the RTE with the Com-Stack) is entirely predetermined by the ECU Extract provided as an input. The required input can be found in the AUTOSAR System Template [14] sections "Data Mapping" and "Communication".

In case the RTE does utilize the Com module for intra-ECU communication it is up to the vendor-specific configuration of the RTE to ensure configuration consistency.

A Metamodel Restrictions

This chapter lists all the restrictions to the AUTOSAR meta-model this version of the AUTOSAR RTE specification document relies on. The RTE generator shall reject configuration where any of the specified restrictions are violated.

A.1 Restriction concerning WaitPoint

1. **[rte_sws_1358]** An error shall be raised if runnable entity has WaitPoint connected to any of the following RTEEvents:
 - OperationInvokedEvent
 - ModeSwitchEvent
 - TimingEvent
 - ExternalEvent
 - DataReceiveErrorEvent

The runnable can only be started with these events.

Rational: For OperationInvokedEvents, ModeSwitchEvents and TimingEvents it suffices to allow the activation of a runnable entity. ExternalEvents are not supported by the RTE SWS of AUTOSAR Release 2.1.

A.2 Restriction concerning RTEEvent

1. **[rte_sws_3526]** The RTE generator shall reject configurations in which a runnable entity which is triggered by the RTEEvent *OperationInvokedEvent* shall be triggered by another RTEEvent too, except if this other RTEEvent is an *OperationInvokedEvent* with compatible operations.

Rationale: The signature of the runnable entity is dependent on its connected RTEEvent.

2. **[rte_sws_3010]** One runnable entity shall only be resumed by one single RTE-Event on its WaitPoint. The RTE doesn't support the WaitPoint of one runnable entity connected to several RTEEvents.

Rationale: The WaitPoint of the runnable entity is caused by calling of the RTE API. One runnable entity can only call one RTE API at a time, and so it can only wait for one RTEEvent.

3. **[rte_sws_7007]** The RTE generator shall reject configurations where RTEEvent instances starting the same runnable entity using implicit data access that are mapped to different OS tasks where one of them might preempt the other.

Rationale: Buffers used for implicit communication shall be consistent during the whole task execution. If it is guaranteed that one task does not preempt the other, direct accesses to the same copy buffer from different tasks are possible.

A.3 Restriction concerning isQueued attribute of DataElement-Prototype

1. **[rte_sws_3012]** Access with DataReadAccess is only allowed for DataElement-Prototypes with their isQueued attribute set to false.

Rationale: By access with DataReadAccess always the last value of the DataElementPrototype will be read in the runnable. There is no meaning to provide a queue of values by DataReadAccess.

2. **[rte_sws_3018]** RTE does not support receiving with WaitPoint for DataElement-Prototypes with their isQueued attribute set to false.

Rationale: "isQueued=false" indicates that the receiver shall not wait for the DataElementPrototype.

3. All the DataSendPoints referring to one DataElementPrototype through one PPort-Prototype are considered to have the same behavior by sending and acknowledgement reception. A DataSendCompletedEvent that references a single DataSendPoint is considered equivalent for all DataSendPoints for the same DataElement-Prototype instance.

Rationale: The API RTESend/RTEWrite is dependent on the port name and the DataElementPrototype name, not on the DataSendPoints. For each combination of one DataElementPrototype and one port only one API will be generated and implemented for sending or acknowledgement reception.

A.4 Restriction concerning ServerCallPoint

1. **[rte_sws_3014]** All the ServerCallPoints referring to one OperationPrototype through one RPortPrototype are considered to have the same behavior by calling service. The RTE generator shall reject configuration where this is violated.

Rationale: The API RTECall is dependent on the port name and the operation name, not on the ServerCallPoints. For each combination of one operation and one port only one API will be generated and implemented for calling a service. It is e.g. not possible to have different timeout values specified for different ServerCallPoints of the same OperationPrototype. It is also not allowed to specify both, a synchronous and an asynchronous server call point for the same OperationPrototype instance.

2. **[rte_sws_3605]** If usage of the indirect API is specified for port prototypes of a software component that all require the same client/server interface, each operation of the client/server interface has to be invoked either by all clients synchronously or by all clients asynchronously.

Rationale: The signature of `Rte_Call` and the existence of `Rte_Result` depend on the kind of invocation.

A.5 Restriction concerning multiple instantiation of software components

1. **[rte_sws_3015]** The RTE only supports multiple objects instantiated from a single AUTOSAR software component by code sharing, the RTE doesn't support code duplication.

Rationale: For AUTOSAR release 2 it was decided to solely concentrate on code sharing and not to support code duplication.

2. **[rte_sws_7101]** The RTE does not support configurations in which a *PortAPI-Option* with *enableTakeAddress = TRUE* is defined by a software-component supporting multiple instantiation.

Rationale: The main focus of the feature for AUTOSAR release 3 was support for configuration of AUTOSAR Services which are limited to single instances.

A.6 Restriction concerning runnable entity

1. **[rte_sws_3016]** The RTE only supports runnable entity of category 1 and 2, the RTE doesn't support runnable entity of category 3.

Rationale: For AUTOSAR release 2 it was decided only to support runnable entity of category 1 and 2, not to support runnable entity of category 3.

2. **[rte_sws_3527]** The RTE does NOT support multiple Runnable Entities sharing the same entry point (symbol attribute of `RunnableEntity`).

Rationale: The handle to data shared by `DataReadAccess` and `DataWriteAccess` has to be coded in the runnable code. An alternative would be an additional parameter to the runnable (a runnable handle) to provide this indirection information.

A.7 Restrictions concerning runnables with dependencies on modes

1. Operations may not be disabled by a `ModeDisablingDependency`.

[rte_sws_2706] RTE shall reject configurations that contain `OperationInvokedEvents` with a `ModeDisablingDependency`.

Rationale: It is a preferable implementation, if the server responds with an explicit application error, when the server operation is not supported in a mode. To implement the disabling of operations would require a high amount of book keeping even for internal client server communication to prevent that the unique request response mapping gets lost.

2. **[rte_sws_2500]** Only a category 1 runnable may be triggered by
 - a `ModeSwitchEvent`
 - an `RteEvent` with a mode disabling dependency

The RTE generator shall reject configurations with category 2 or 3 runnables connected to `ModeSwitchEvents` and `RteEvents` with mode disabling dependencies.

Rationale: The above runnables are executed or terminated on the transitions between different modes. To execute the mode switch withing finite time, also these runnables have to be executed within finite execution time.

3. All `OnEntry` runnables and `OnExit` runnables of the same mode machine instance should be mapped to the same task.

[rte_sws_2662] The RTE generator shall reject configurations with `OnEntry` or `OnExit` runnables of the same mode machine instance that are mapped to different tasks.

Rationale: This restriction simplifies the implementation of the semantics of a mode switch.

4. To guarantee that all mode disabling dependent runnables of a mode machine instance have terminated before before the start of the `OnExit` runnables of the transition, the mode disabling dependent runnables should run with higher or equal priority.

[rte_sws_2663] The RTE generator shall reject configurations with mode disabling dependent runnables that are mapped to a task with lower priority than the task that contains the `OnEntry` runnables and `OnExit` runnables of that mode machine instance.

5. **[rte_sws_2664]** The RTE generator shall reject configurations of a task with `OnExit` runnables mapped behind `OnEntry` runnables of the same mode machine instance.

Rationale: This restriction simplifies the implementation of the semantics of a mode switch.

6. If a mode is used to trigger a runnable for entering or leaving the mode, but this runnable has a mode disabling dependency on the same mode, the mode

disabling dependency inhibits the activation of the runnable on the transition (see section 4.4.4).

To prevent such a misleading configuration, it is strongly recommended not to configure a mode disabling dependency for an `OnEntry runnable` or `OnExit runnable`, using the same mode.

A.8 Restriction concerning InterRunnableVariables

1. **[rte_sws_3518]** The usage of *InterRunnableVariables with implicit behavior* shall be valid for category 1a and 1b Runnable entities only.

Rationale: The update of *InterRunnableVariables with implicit behavior* done during a Runnable execution shall be made available to other Runnables after the Runnable execution has terminated (see `rte_sws_3584`). This limitation is not valid for *InterRunnableVariables with explicit behavior*.

Runnable termination is not guaranteed for Runnables of category 2 or 3.

2. **[rte_sws_3588]** *InterRunnableVariables* don't support complex data types.

Rationale: If *InterRunnableVariables* would support complex data types, a reference would have to be passed for read access. Afterwards Runnable code will access the complex data type via the reference. But RTE is only able to protect (for data consistency purposes) the delivery of the reference, not the access to the referenced data later on. In those cases, when complex data has to be used for Intra AUTOSAR SW-C communication it must be sufficient to apply `ExclusiveAreas` (see section 4.2.4.5 and API in section 5.6.24 and 5.6.25) to force the RTE guaranteeing data consistency.

3. **[rte_sws_3591]** *InterRunnableVariables* don't support the AUTOSAR primitive type string

Rationale: In those cases when a string is used for Intra AUTOSAR SW-C communication it should be sufficient to apply `ExclusiveAreas` (see section 4.2.4.5 and API in section 5.6.24 and 5.6.25) to force the RTE guaranteeing data consistency.

The mass of *InterRunnableVariables* is expected to be of other primitive types than strings. Support for strings might be added in a later release together with support of complex data. Both require reference passing. Also see `rte_sws_3588`.

A.9 Restriction concerning InternalBehavior

1. **[rte_sws_5034]** There shall only be one *InternalBehavior* provided for each *AtomicSoftwareComponentType*.

Rationale: For the generation of the application header file not only the *AtomicSoftwareComponentType* but also the *InternalBehavior* is relevant. In case two implementation for the same *AtomicSoftwareComponentType* – but different *InternalBehavior* – are mapped to the same ECU two application header files for the same *AtomicSoftwareComponentType* would be required. In this document release the application header file is defined based on the *AtomicSoftwareComponentType*, therefore it is not allowed to specify different *InternalBehavior* additionally.

In a future release the application header file shall be based on the *InternalBehavior*, then this restriction is not valid anymore.

A.10 Restriction concerning Initial Value

1. [rte_sws_4525] Each instance within one ECU of a data element that is connected to the same sender is imposed to use identical init values.

Rationale: In the meta model init values are specified in the data receiver com spec. Since a separate data receiver com spec exists for each port that categorizes a specific interface, it would be (theoretically) possible to define a different init value for a certain data element in each port. But COM allows only one init value per signal.

A.11 Restriction concerning PerInstanceMemory

1. [rte_sws_3790] The `<typeDefinition>` attribute of a `PerInstanceMemory` is not allowed to contain a function pointer.

Rationale: Using the typedefinition `typedef <typedefinition> <typename>` does not work for function pointers.

A.12 Restriction concerning unconnected r-port

1. [rte_sws_3019] If strict checking has been enabled (see `rte_sws_5099`) there shall not be unconnected r-port. The RTE generator shall in this case reject the configuration with unconnected r-port.

Rationale: Unconnected r-port is considered as wrong configuration of the system.

1. [rte_sws_2750] The RTE Generator shall reject configurations where an r-port typed with a `CalprmInterface` is not connected and an `initValue` of a `ParameterRequireComSpec` is not provided for each `CalprmElementPrototype` of this `CalprmInterface`.

A.13 Restrictions regarding n:1 sender-receiver communication

1. [rte_sws_2670] RTE shall not support connections with multiple senders (n:1 communication) of `mode switch notifications` connected to the same receiver. The RTE generator shall reject configurations with multiple senders of `mode switch notifications` connected to the same receiver.

Rationale: No use case is known to justify the required complexity.

A.14 Restrictions regarding Measurement and Calibration

1. [rte_sws_3951] RTE does not support measurement of queued communication.

Rationale: Measurement of queued communication is not supported yet. Reasons are:

- A queue can be empty. What's to measure then? Data interpretation is ambiguous.
 - Which of the queue entries the measurement data has to be taken from (first pending entry, last entry, an intermediate one, mean value, min. or max. value)? Needs might differ out of user view? Data interpretation is ambiguous.
 - Compared e.g. to sender-receiver last-is-best approach only inefficient solutions are possible because implementation of queues entails storage of information dynamically at different memory locations. So always additional copies are required.
2. [rte_sws_3970] The RTE generator shall reject configurations containing require ports attached to `CalprmComponentTypes`.

Rationale: Require ports on `CalprmComponentTypes` don't make sense. `CalprmComponentTypes` only have to provide calibration parameters to other Component types.

A.15 Restriction concerning `ExclusiveAreaImplMechanism`

1. [rte_sws_3811] If an exclusive area's configuration value for `ExclusiveAreaImplMechanism` is `InterruptBlocking`, `OsResource`, or `NonPreemptiveTasks`, no runnable entity shall contain any waitpoint inside this exclusive area.

Please note that a waitpoint can either be a modelling waitpoint e. g. a waitpoint in the SW-C description caused by the usage of a blocking API (e. g. `Rte_Receive`) or an implementation waitpoint caused by a special implementation to fulfill the requirements of the ECU configuration, e. g. the runnable-to-task mapping.

Rationale: The operating system has the limitation that a `WaitEvent` call is not allowed with disabled interrupts. Therefore the implementation mechanism *InterruptBlocking* cannot be used if the exclusive area contains a waitpoint.

Further the operating system has the limitation that an OS waitpoint cannot be entered with occupied OS Resources. This implies that the implementation mechanism *OsResource* cannot be used if the exclusive area contains a waitpoint.

A runnable entity containing a waitpoint cannot be mapped to a non-preemptive task, because the waitpoint is a rescheduling point. Therefore the implementation mechanism *NonPreemptiveTasks* cannot be used if the runnable entity using the exclusive area contains a waitpoint.

B Required Input Information

This chapter lists all the input information necessary for the RTE generator in a tabular form. The meanings of the individual field entries are described in the following table:

Requirement ID	Unique ID of the RTE SWS input requirement.
Object identifier	Unique identifier in the RTE SWS representing the metamodel object. It is used like a variable in the RTE SWS standing for the content of an instance of the described metamodel object.
Object information	Necessary object information required for RTE generation in terms of a short description.
Description	Description of the required object information in more detail. It may contain a listing of the possible values of the required input information and constraints.
Rationale	Reason why the described metamodel object is needed as an input to the RTE generation.
Template metamodel path	Metamodel path of the object in an AUTOSAR template, e. g. "AUTOSAR Software Component Template"[17] or "System Template" [14].
Required by	Lists all RTE SWS requirements that depend on the existence of the described metamodel object.
Contract phase	Specifies whether the input information is already required for the contract phase.

"M2" in the template metamodel path means "metamodel" level (see "AUTOSAR Template Modeling Guide" [16]). This document is especially important to understand the specific semantics of the AUTOSAR metamodel (like the semantics of the "instanceRef" and "isOfType" associations).

In certain cases, some attributes of a class are not given directly in the class-table, if they are inherited from the base classes. For example, attribute "Identifier" of class "ComponentType" is not included in the class-table, because it is inherited from the base class "ARElement", which again inherits the attribute from its base class "Identifiable".

B.1 SWC and instance

Requirement ID	[rte.sws.in.0001]
Object identifier	SwcTypeName
Object information	Name of each SWC type
Description	Defines the name of the software component type. Shall be unique within the ECU.
Rationale	To define the API mapping in the Application Header File. Define the Component Data Structure in the generated RTE.

Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::Components::ComponentType::Identifier
Required by	rte_sws_1003 rte_sws_1143 rte_sws_1155 rte_sws_1348 rte_sws_3714 rte_sws_3731
Contract phase	Yes

Requirement ID	[rte_sws_in_0002]
Object identifier	SwcImplementationLanguage
Object information	Implementation language of each SWC
Description	For the implementation language of software components currently only C/C++ are supported.
Rationale	To define the using of C linkage in the Application Header File
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::Implementation::programmingLanguage
Required by	rte_sws_1011
Contract phase	No

Requirement ID	[rte_sws_in_0003]
Object identifier	SwcSourceCodeDelivery
Object information	Source Code availability of the SWC
Description	Whether or not the source code is available for a SWC
Rationale	To decide if the Application Header File can be optimized again by RTE-Gen phase.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::Implementation::Code::type
Required by	rte_sws_1216
Contract phase	No

Requirement ID	[rte_sws_in_0004]
Object identifier	supportsMultipleInstantiation
Object information	Multi-Instantiation of the SWC
Description	Whether the SWC can be multiply instantiated
Rationale	To define the API mapping in the Application Header File.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::InternalBehavior::supportsMultipleInstantiation
Required by	rte_sws_2008 rte_sws_2009 rte_sws_3706 rte_sws_3707
Contract phase	Yes

Requirement ID	[rte_sws_in_0009]
----------------	--------------------------

Object identifier	PerInstanceMemoryName
Object information	Name of each PerInstanceMemory when attribute supportsMultipleInstantiation==TRUE
Description	The name of a PerInstanceMemory shall be unique within the SWC.
Rationale	To define the name of the PerInstanceMemory handle and the API mapping in the Application Header File and allocate the PerInstanceMemory in the generated rte.c.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::InternalBehavior::PerInstanceMemory::Identifier
Required by	rte_sws_1118 rte_sws_2305 rte_sws_2301
Contract phase	Yes

Requirement ID	[rte_sws_in_0071]
Object identifier	PerInstanceMemoryType
Object information	Name of the type of each PerInstanceMemory when attribute supportsMultipleInstantiation==TRUE
Description	The type name of a PerInstanceMemory.
Rationale	To define the type of the PerInstanceMemory handle in the Application Header File and allocate the PerInstanceMemory in the generated rte.c.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::InternalBehavior::PerInstanceMemory::type
Required by	rte_sws_1118 rte_sws_2303 rte_sws_2302
Contract phase	Yes

Requirement ID	[rte_sws_in_0068]
Object identifier	PerInstanceMemoryTypeDef
Object information	Type definition of each PerInstanceMemory when attribute supportsMultipleInstantiation==TRUE
Description	The type definition of a PerInstanceMemory shall be in valid c-syntax.
Rationale	To define the type of the PerInstanceMemory in the Application Header File.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::InternalBehavior::PerInstanceMemory::typeDefinition
Required by	rte_sws_1118 rte_sws_2304
Contract phase	Yes

Requirement ID	[rte_sws_in_5061]
Object identifier	RamBlockLocationSymbol
Object information	Name of the PerInstanceMemory symbol name to be generated
Description	When instantiating the PerInstanceMemory for the usage as a RamBlock for the NvRam Manager the RTE generator shall use this specified name. The name has to be unique for the whole ECU.

Rationale	The name of the PerInstanceMemory instance has to be available for the configuration of the NvRam Manager.
Template meta-model path	M1::AUTOSAR Descriptions::ECUCParameterDefinition::RTE::NVRamAllocation::RamBlockLocationSymbol
Required by	rte_sws_5062
Contract phase	No

Requirement ID	[rte_sws_in_3750]
Object identifier	RequiredRteOperatingMode
Object information	Required RTE Operating Mode
Description	An AUTOSAR software component shall indicate its required operating mode.
Rationale	Based on this attribute the RTE Generator can perform optimizations.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::InternalBehavior::Implementation::RTEVendor
Required by	rte_sws_1234
Contract phase	No

Requirement ID	[rte_sws_in_5013]
Object identifier	Constants
Object information	Published Constants
Description	Each constant defined in the SW-Component description will be accessed and published.
Rationale	The Application Header File shall make visible the constants encountered in the input using the appropriate AUTOSAR data-types.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::Datatype::Constants::Constant
Required by	
Contract phase	Yes

Requirement ID	[rte_sws_in_5046]
Object identifier	EcuAbstractionSWComponent
Object information	Reference to the SW-Component which represents the EcuAbstraction
Description	With this reference to the local EcuAbstraction it is possible to distinguish between the EcuAbstraction and other kinds of SW-Components.
Rationale	The RTE needs to make sure that no communication to an remote EcuAbstraction is configured. Therefore a reference to the local EcuAbstraction needs to be provided.
Template meta-model path	

Required by	rte_sws_2051
Contract phase	No

B.2 Runnable entity and task

Requirement ID	[rte_sws_in_0012]
Object identifier	RunnableEntityName
Object information	Name of each runnable entity
Description	Shall be unique within the SWC
Rationale	To define the API in the Application Header File.
Template meta-model path	M2::AUTOSAR_Templates::SWComponentTemplate::InternalBehavior::RunnableEntity::Identifier
Required by	rte_sws_3733 rte_sws_3741 rte_sws_3744
Contract phase	Yes

Requirement ID	[rte_sws_in_0013]
Object identifier	RunnableEntityToTaskMapping
Object information	Mapping of runnables to OS tasks
Description	Defines the mapping of the Runnable Entity instances to OS Tasks.
Rationale	Generate the task body content.
Template meta-model path	M1::AUTOSAR_Descriptions::ECUCParameterDefinition::RTE::Tasks::RunnableEntityMapping
Required by	rte_sws_2204 rte_sws_2251
Contract phase	No

Requirement ID	[rte_sws_in_5012]
Object identifier	TaskBodyName
Object information	Name of the generated task body
Description	The names of the generated task bodies have to be unique on one ECU. The name is the shortName of the corresponding OsTask.
Rationale	Generate the C module containing the task body.
Template meta-model path	M1::AUTOSAR_Descriptions::ECUCParameterDefinition::Services::OS::OsTask::shortName
Required by	rte_sws_1257 rte_sws_2251 rte_sws_4014
Contract phase	No

Requirement ID	[rte_sws_in_0040]
Object identifier	OObjects

Object information	ECU configuration parameters of the AUTOSAR OS
Description	The RTE generator needs access to the ECU-Configuration parameters of the AUTOSAR OS.
Rationale	Determine the type of a task
Template meta-model path	M1::AUTOSAR Descriptions::ECUCParameterDefinition::Services::OS::OsTask
Required by	rte.sws.2251 rte.sws.4014
Contract phase	No

Requirement ID	[rte.sws.in.0014]
Object identifier	RunnableEntitySequence
Object information	Sequences of Runnable Entities in each OS task
Description	Defines the sequence the Runnable Entities are called within one task body.
Rationale	Generate the task body content.
Template meta-model path	M1::AUTOSAR Descriptions::ECUCParameterDefinition::RTE::Tasks::PositionInTask
Required by	rte.sws.2207
Contract phase	No

Requirement ID	[rte.sws.in.0053]
Object identifier	EntryPointSymbol
Object information	Symbol describing a runnable's entry point
Description	A runnable is represented as a function in C/C++ code. This symbol represents the entry point of the function.
Rationale	The entry point symbol is considered to be the API name of the runnable.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::InternalBehavior::RunnableEntity::symbol
Required by	rte.sws.1126 rte.sws.1131 rte.sws.1133 rte.sws.1135 rte.sws.1137 rte.sws.1166 rte.sws.2512
Contract phase	Yes

Requirement ID	[rte.sws.in.0015]
Object identifier	OsTaskPriority
Object information	Priority of each OS task
Description	Provide the priority of each OS Task.

Rationale	The ECU Configuration has to ensure that a server runnable of a synchronous C/S call that cannot be invoked as a direct function call is mapped to a task with a higher priority than the calling client runnable
Template meta-model path	M1::AUTOSAR Descriptions::ECUCParameterDefinition::Services::OS::OsTask::OsTaskPriority
Required by	rte_sws_2251 rte_sws_4014
Contract phase	No

Requirement ID	[rte_sws_in_5070]
Object identifier	ActivationOffset
Object information	Activation Offset in seconds
Description	Provides the input which activation offset shall be configured for this runnable entity mapping.
Rationale	The RTE shall respect the configured activation offset of runnable entities mapped within one OS task.
Template meta-model path	M1::AUTOSAR Descriptions::ECUCParameterDefinition::RTE::RunnableTaskMapping::ActivationOffset
Required by	rte_sws_7000
Contract phase	No

Requirement ID	[rte_sws_in_0039]
Object identifier	OsEvent
Object information	Name of the OSEvent
Description	The OSEvent to which the RTEEvent is assigned
Rationale	For the RTEEvents which are implemented with OSEvents the name of the OSEvents shall be defined.
Template meta-model path	M1::AUTOSAR Descriptions::ECUCParameterDefinition::RTE::Tasks::RunnableEntityMapping::UsedOsEventRef
Required by	rte_sws_2251 rte_sws_4014
Contract phase	No

Requirement ID	[rte_sws_in_5016]
Object identifier	ExclusiveAreaName
Object information	Name of the exclusive area
Description	The Internal Behavior does provide the list of defined exclusive areas.
Rationale	Define the name of the handle for the exclusive area.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::InternalBehavior::ExclusiveArea::ExclusiveArea::Identifier
Required by	rte_sws_3739
Contract phase	No

Requirement ID	[rte_sws_in_5017]
Object identifier	InterRunnableVariableName
Object information	Name of the Interrunnable Variable
Description	The Internal Behavior does provide this list of defined inter runnable variables.
Rationale	Generate the Application Header File for exclusive are access.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::InternalBehavior::InterRunnableVariable::Identifier
Required by	rte_sws_1120 rte_sws_1123 rte_sws_3550 rte_sws_3553 rte_sws_3560 rte_sws_3565
Contract phase	Yes

Requirement ID	[rte_sws_in_0070]
Object identifier	RTEEvent
Object information	RTE Event
Description	The RTE Event which triggers the runnable entity
Rationale	Define the trigger conditions of the runnable entities
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::InternalBehavior::RTEEvents
Required by	rte_sws_2203
Contract phase	Yes

Requirement ID	[rte_sws_in_0072]
Object identifier	RunnableEntityInvokedConcurrently
Object information	the attribute canBelInvokedConcurrently of Runnable Entity
Description	whether the runnable entity can be invoked concurrently
Rationale	for task mapping
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::InternalBehavior::RunnableEntity::canBelInvokedConcurrently
Required by	rte_sws_3523
Contract phase	Yes

Requirement ID	[rte_sws_in_0073]
Object identifier	DataReadAccess
Object information	the attribute dataReadAccess of Runnable Entity
Description	the implicit read access of a RunnableEntity to a DataElement
Rationale	Defines the data read access behavior of a RunnableEntity. It is necessary for API definition.

Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::InternalBehavior::RunnableEntity::dataReadAccess
Required by	rte_sws_6000 rte_sws_6001 rte_sws_6004 rte_sws_6011
Contract phase	Yes

Requirement ID	[rte_sws_in_0074]
Object identifier	DataReceivePoint
Object information	the attribute dataReceivePoint of Runnable Entity
Description	the explicit read access of a RunnableEntity to a DataElement
Rationale	Defines the data read access behavior of a RunnableEntity. It is necessary for API definition.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::InternalBehavior::RunnableEntity::dataReceivePoint
Required by	rte_sws_6011
Contract phase	Yes

Requirement ID	[rte_sws_in_0075]
Object identifier	DataSendPoint
Object information	the attribute dataSendPoint of Runnable Entity
Description	the explicit write access of a RunnableEntity to a DataElement
Rationale	Defines the data write access behavior of a RunnableEntity. It is necessary for API definition.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::InternalBehavior::RunnableEntity::dataSendPoint
Required by	rte_sws_6011 rte_sws_6016
Contract phase	Yes

Requirement ID	[rte_sws_in_0076]
Object identifier	DataWriteAccess
Object information	the attribute dataWriteAccess of Runnable Entity
Description	the implicit write access of a RunnableEntity to a DataElement
Rationale	Defines the data write access behavior of a RunnableEntity. It is necessary for API definition.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::InternalBehavior::RunnableEntity::dataWriteAccess
Required by	rte_sws_6011 rte_sws_3570 rte_sws_3571
Contract phase	Yes

Requirement ID	[rte_sws_in_0079]
----------------	--------------------------

Object identifier	ServerCallPoint
Object information	the attribute serverCallPoint of Runnable Entity
Description	The RunnableEntity has a serverCallPoint to the referenced operation
Rationale	References the operation that the RunnableEntity can call.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::InternalBehavior::RunnableEntity::serverCallPoint
Required by	rte_sws_1293 rte_sws_1294
Contract phase	No

Requirement ID	[rte_sws_in_0081]
Object identifier	WaitPoint
Object information	the attribute waitPoint of Runnable Entity
Description	The RunnableEntity has a waitPoint to the referenced RTEEvent
Rationale	References the RTEEvent that the RunnableEntity can wait for.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::InternalBehavior::RunnableEntity::waitPoint
Required by	rte_sws_1290
Contract phase	No

Requirement ID	[rte_sws_in_2700]
Object identifier	minimum start interval
Object information	minimum start interval of a runnable
Description	the minimum start interval of a runnable gives the minimum interval between two starts of a runnable. minimum start interval 0: no start interval monitoring
Rationale	Required to prevent a high activation frequency
Template meta-model path	tbd
Required by	rte_sws_2697
Contract phase	No

B.3 Port and interface

Requirement ID	[rte_sws_in_0018]
Object identifier	PortName
Object information	Name of the port
Description	Shall be unique within the SWC
Rationale	To identify different port prototype for API generation

Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::Components::PortPrototype::Identifier
Required by	rte_sws_1071 rte_sws_1072 rte_sws_1206 rte_sws_1083 rte_sws_1091 rte_sws_1092 rte_sws_1102 rte_sws_1111 rte_sws_3741 rte_sws_3744
Contract phase	Yes

Requirement ID	[rte_sws_in_0019]
Object identifier	RPort/PPort
Object information	Type of the port
Description	r- or p- port
Rationale	To indicate whether the port is provided or required port for configuration checking
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::Components::PortPrototype
Required by	rte_sws_5508
Contract phase	Yes

Requirement ID	[rte_sws_in_1352]
Object identifier	InterfaceName
Object information	Name of the interface
Description	Shall be unique within the system
Rationale	To ensure unique names for those things that are related to a particular interface rather than the ports that are characterized by the interface
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::PortInterface::Identifier
Required by	rte_sws_2576
Contract phase	Yes

Requirement ID	[rte_sws_in_0069]
Object identifier	InterfacelsService
Object information	isService attribute of the PortInterface
Description	Whether port provides or requires the interface is a service port
Rationale	To distinguish the communication with normal SWC and the communication with Basic-SW services.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::PortInterface::isService
Required by	rte_sws_2100
Contract phase	Yes

Requirement ID	[rte_sws_in_0020]
Object identifier	DataElementName
Object information	Name of the data element
Description	Shall be unique within the SWC
Rationale	To identify different data element prototype for API generation
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::PortInterface::DataElementPrototype::Identifier
Required by	rte_sws_1071 rte_sws_1072 rte_sws_1206 rte_sws_1083 rte_sws_1091 rte_sws_1092 rte_sws_3741 rte_sws_3744
Contract phase	Yes

Requirement ID	[rte_sws_in_0060]
Object identifier	DataElementDatatype
Object information	Data type of the data element
Description	Contains the information like upper/lower-limit for integer and real type
Rationale	For API generation
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::Datatype::Datatypes
Required by	rte_sws_1071 rte_sws_1072 rte_sws_1206 rte_sws_1083 rte_sws_1091 rte_sws_1092 rte_sws_3741 rte_sws_3744
Contract phase	Yes

Requirement ID	[rte_sws_in_0045]
Object identifier	DataElementIsQueued
Object information	Specifies whether the data element is queued or not. VFB attribute: INFORMATION_TYPE
Description	Qualifies whether the content of the data element is queued. If it is queued then the data element has event semantics - i.e. data elements are stored in a queue and all data elements are processed in first in first out order. If it is not queued then the last is best semantics applies.
Rationale	For configuration checking and API generation
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::PortInterface::DataElementPrototype::isQueued
Required by	rte_sws_1071 rte_sws_1072 rte_sws_5033
Contract phase	Yes

Requirement ID	[rte_sws_in_0058]
Object identifier	OperationName
Object information	Name of the operation
Description	Shall be unique within the SWC

Rationale	To identify different operation prototype for API generation
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::PortInterface::OperationPrototype::Identifier
Required by	rte_sws_1102 rte_sws_1111
Contract phase	Yes

Requirement ID	[rte_sws_in_0059]
Object identifier	ArgumentName
Object information	Name of the argument of the operation
Description	Shall be unique within the operation
Rationale	For API generation
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::PortInterface::OperationPrototype::ArgumentPrototype::Identifier
Required by	rte_sws_1102 rte_sws_1111
Contract phase	Yes

Requirement ID	[rte_sws_in_0061]
Object identifier	ArgumentDirection
Object information	Direction of the argument of the operation
Description	In/Out/Inout
Rationale	For API generation
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::PortInterface::OperationPrototype::ArgumentPrototype::Direction
Required by	rte_sws_1102 rte_sws_1111
Contract phase	Yes

Requirement ID	[rte_sws_in_0021]
Object identifier	AssemblyConnectorPrototype
Object information	Connection of communication partners (ports)
Description	Refers to one p-port and one r-port
Rationale	For API implementation
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::Composition::AssemblyConnectorPrototype
Required by	rte_sws_2200
Contract phase	No

Requirement ID	[rte_sws_in_0055]
Object identifier	SInitValue

Object information	Initial value of a data element prototype (isQueued = false) on the sender side. VFB attribute on sender side: INIT_VALUE.
Description	Refers to a constant value. Only the ComSpec of an AtomicSoftware-ComponentType PortPrototype shall be considered.
Rationale	To prevent calculation based on invalid values
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::Communication::DataSenderComSpec::InitValue
Required by	rte.sws.6009 rte.sws.6010
Contract phase	No

Requirement ID	[rte.sws.in.0062]
Object identifier	RInitValue
Object information	Initial value of a data element prototype (isQueued = false) on the sender side. VFB attribute on receiver side: INIT_VALUE
Description	Refers to a constant value. Only the ComSpec of an AtomicSoftware-ComponentType PortPrototype shall be considered.
Rationale	To prevent calculation based on invalid values
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::Communication::DataReceiverComSpec::InitValue
Required by	rte.sws.6010
Contract phase	No

Requirement ID	[rte.sws.in.0023]
Object identifier	ServerRunnable
Object information	for each operation the connected runnable entity
Description	Refers to the runnable entity which shall be activated when the OperationInvokedEvent is triggered
Rationale	For invocation of the server runnable
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::InternalBehavior::RTEEvents::OperationInvokedEvent::RunnableEntityRef
Required by	rte.sws.1166
Contract phase	No

Requirement ID	[rte.sws.in.2574]
Object identifier	ApplicationErrorValues
Object information	Application Error Value definition for each operation
Description	The definition of the Application Error Values used in exchange between SW-Components (with symbolic name and value)
Rationale	Application Errors shall be defined in the Application Header File. For definition of Rte_StatusType.

Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::VFBErrors::ApplicationError
Required by	rte_sws_2573 rte_sws_2575 rte_sws_2576
Contract phase	Yes

Requirement ID	[rte_sws_in_5023]
Object identifier	CanInvalidate
Object information	Can the sender invalidate the data element
Description	When specified the sender of a data element can set the value to the invalid value defined in the data semantics. Only the ComSpec of an AtomicSoftwareComponentType PortPrototype shall be considered.
Rationale	For API generation of data element invalidation
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::Communication::DataSenderComSpec::canInvalidate
Required by	rte_sws_5024
Contract phase	Yes

Requirement ID	[rte_sws_in_5031]
Object identifier	InvalidValue
Object information	Invalid value
Description	The value to be used when invalidating a data element.
Rationale	The value to be used for the invalid data indication must be the same for all partners in the communication.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::Datatype::Datatypes::PrimitiveType::SwDataDefProps::invalidValue
Required by	rte_sws_3802 rte_sws_5025
Contract phase	Yes

Requirement ID	[rte_sws_in_5050]
Object identifier	handleInvalid
Object information	handleInvalid [keep ; replace]
Description	Specifies at the UnqueuedReceiverComSpec whether an received invalid value shall be kept or replaced. Only the ComSpec of an AtomicSoftwareComponentType PortPrototype shall be considered.
Rationale	The receiver RTE / COM needs to be configures what to do when an invalid value is received.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::Communication::UnqueuedReceiverComSpec::handleInvalid
Required by	rte_sws_5032 rte_sws_5026 rte_sws_5048 rte_sws_5030 rte_sws_5049
Contract phase	No

Requirement ID	[rte_sws_in_3777]
Object identifier	TransmissionAcknowledgementRequest
Object information	Request a transmission acknowledgment
Description	Requests acknowledgments that data has been sent successfully.
Rationale	The sender of a data element can request an acknowledgment for successful or erroneous transmission using this attribute
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::Communication::TransmissionAcknowledgementRequest
Required by	rte_sws_5504 rte_sws_5506 rte_sws_3754 rte_sws_3755
Contract phase	Yes

Requirement ID	[rte_sws_in_1361]
Object identifier	PortDefinedArgumentType
Object information	Data type of port-defined argument
Description	The data type that the server runnable entity requires to be passed.
Rationale	To enable correct function prototypes to be emitted
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::InternalBehavior::PortAPIOptions::type
Required by	rte_sws_1166
Contract phase	Yes

Requirement ID	[rte_sws_in_1362]
Object identifier	PortDefinedArgumentValue
Object information	Value of port-defined argument
Description	Value to pass for a specific port-defined argument for a specific server SWC (instance).
Rationale	To enable correct values to be passed as the port-defined arguments for invocation of server runnables.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::Internal-Behavior::PortAPIOptions::value
Required by	rte_sws_1360
Contract phase	No

Requirement ID	[rte_sws_in_3798]
Object identifier	indirectAPI
Object information	Selection of indirect API
Description	If indirectAPI = true the indirect API shall be generated for the referenced port prototype.

Rationale	To avoid generating unnecessary entries of the component data structure.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::Internal-Behavior::PortAPIOptions::indirectAPI
Required by	rte_sws_3799 rte_sws_3601 rte_sws_2613 rte_sws_2615 rte_sws_1355
Contract phase	Yes

Requirement ID	[rte_sws_in_7102]
Object identifier	enableTakeAddress
Object information	API reference for derivation of a pointer
Description	If enableTakeAddress = true the software-component is able to use the API reference for deriving a pointer to an object.
Rationale	To guarantee that a pointer to an object for the port API can be derived if required by the implementation of a software-component.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::Internal-Behavior::PortAPIOptions::enableTakeAddress
Required by	rte_sws_7100 rte_sws_7101
Contract phase	Yes

Requirement ID	[rte_sws_in_1361]
Object identifier	PortDefinedArgumentType
Object information	Data type of port-defined argument
Description	The data type that the server runnable entity requires to be passed.
Rationale	To enable correct function prototypes to be emitted
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::InternalBehavior::PortArgument::type
Required by	rte_sws_1166
Contract phase	Yes

Requirement ID	[rte_sws_in_1362]
Object identifier	PortDefinedArgumentValue
Object information	Value of port-defined argument
Description	Value to pass for a specific port-defined argument for a specific server SWC (instance).
Rationale	To enable correct values to be passed as the port-defined arguments for invocation of server runnables.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::Internal-Behavior::PortArgument::value
Required by	rte_sws_1360
Contract phase	No

B.4 Communication

Requirement ID	[rte_sws_in_0067]
Object identifier	AliveTimeout
Object information	The minimum time period for the reception of the data element (isQueued = false). VFB attribute: LIVELIHOOD
Description	When specified the receiver can monitor the time-out and inform a time-out to the software component. Only the ComSpec of an AtomicSoftwareComponentType PortPrototype shall be considered.
Rationale	For API generation of the time-out notification
Template meta-model path	M2::AUTOSAR_Templates::SWComponentTemplate::Communication::DataReceiverComSpec::aliveTimeout
Required by	rte_sws_5020 rte_sws_5021 rte_sws_5022
Contract phase	Yes

Requirement ID	[rte_sws_in_0066]
Object identifier	RFiltering
Object information	The filter mechanism on the receiver side. SWCT attribute: filter
Description	of class DataFilter
Rationale	For API implementation to filter the data element according to certain mechanism on the receiver side
Template meta-model path	M2::AUTOSAR_Templates::SWComponentTemplate::Communication::Filter::DataFilter
Required by	rte_sws_5503
Contract phase	No

Requirement ID	[rte_sws_in_0029]
Object identifier	QueuedReceiverComSpec.QueueLength
Object information	The length of the queue of the received data element (isQueued = true)
Description	of type Integer. Only the ComSpec of an AtomicSoftwareComponentType PortPrototype shall be considered.
Rationale	For configuration of the queue
Template meta-model path	M2::AUTOSAR_Templates::SWComponentTemplate::Communication::QueuedReceiverComSpec::QueueLength
Required by	rte_sws_2521
Contract phase	No

Requirement ID	[rte_sws_in_2701]
Object identifier	ServerComSpec.QueueLength
Object information	The length of the queue of requests to a serialised server operation

Description	of type Integer. Only the ComSpec of an AtomicSoftwareComponent-Type PortPrototype shall be considered.
Rationale	For configuration of the queue
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::Communication::ServerComSpec::queueLength
Required by	rte_sws_2529 rte_sws_2530 rte_sws_2699
Contract phase	No

Requirement ID	[rte_sws_in_0063]
Object identifier	SignalMappingP
Object information	Mapping of primitive data element to COM signal(s)
Description	refers to data element instance and the COM signal(s) - the COM signal is the interface of COM to RTE.
Rationale	For API implementation by invocation of COM API
Template meta-model path	M2::AUTOSAR Templates::SystemTemplate::DataMapping::SenderReceiverToSignalMapping
Required by	rte_sws_3007 rte_sws_4504 rte_sws_4505
Contract phase	No

Requirement ID	[rte_sws_in_0064]
Object identifier	SignalMappingC
Object information	Mapping of complex data element to COM signal group(s)
Description	refers to data element instance and the COM signal group(s) - the COM signal group is the interface of COM to RTE.
Rationale	For API implementation by invocation of COM API
Template meta-model path	M2::AUTOSAR Templates::SystemTemplate::DataMapping::SenderReceiverToSignalGroupMapping
Required by	rte_sws_3008 rte_sws_4506 rte_sws_4507 rte_sws_4508 rte_sws_2557
Contract phase	No

Requirement ID	[rte_sws_in_5068]
Object identifier	OperationMapping
Object information	Mapping of ClientServer operations to SignalGroups
Description	Specifies how the individual operation arguments are mapped on SystemSignals and SystemSignalGroups.
Rationale	For API implementation by invocation of COM API
Template meta-model path	M2::AUTOSAR Templates::SystemTemplate::DataMapping::ClientServerToSignalGroupMapping
Required by	rte_sws_5055
Contract phase	No

Requirement ID	[rte_sws_in_5069]
Object identifier	ClientServerProtocolMapping
Object information	Mapping of the ClientServer call/response to SignalGroups
Description	Specifies how the client server call/response is mapped on SystemSignals and SystemSignalGroups.
Rationale	For API implementation by invocation of COM API
Template meta-model path	M2::AUTOSAR Templates::SystemTemplate::DataMapping::ClientServerToSignalGroupMapping
Required by	rte_sws_5054 rte_sws_5055 rte_sws_6028 rte_sws_5056 rte_sws_5057 rte_sws_5058 rte_sws_5059
Contract phase	No

Requirement ID	[rte_sws_in_5079]
Object identifier	ComSignalHandleId
Object information	Reference from ComSignal to signal instance.
Description	Reference to the ISignalToIPduMapping that contains an ISignal (System Template) which this ComSignal represents.
Rationale	To extract the HandleId which needs to be used for interaction between Rte and Com
Template meta-model path	M1::AUTOSAR Descriptions::ECUCParameterDefinition::COM-Stack::Com::SystemTemplateSystemSignalRef
Required by	rte_sws_4504 rte_sws_4505
Contract phase	No

Requirement ID	[rte_sws_in_5080]
Object identifier	ComSignalGroupHandleId
Object information	Reference from ComSignalGroup to signal instance.
Description	Reference to the ISignalToIPduMapping that contains an ISignal (SystemTemplate) which this ComSignalGroup represents.
Rationale	To extract the HandleId which needs to be used for interaction between Rte and Com
Template meta-model path	M1::AUTOSAR Descriptions::ECUCParameterDefinition::COM-Stack::Com::SystemTemplateSystemSignalRef
Required by	rte_sws_4506 rte_sws_4507 rte_sws_4508 rte_sws_2557 rte_sws_5081
Contract phase	No

B.5 Data consistency

Requirement ID	[rte_sws_in_3597]
Object identifier	ExclusiveAreaImplMechanism

Object information	ExclusiveArea data consistency mechanism
Description	Parameter specifying the data consistency mechanism to be applied to an ExclusiveArea
Rationale	Influence RTE behavior allowing specific optimizations in view of usage of ECU resources
Template meta-model path	M1::AUTOSAR Descriptions::ECUCParameterDefinition::RTE::Data-Consistency::ExclusiveAreaImplMechanism
Required by	rte_sws_3503 rte_sws_3504 rte_sws_3595 rte_sws_3596
Contract phase	No

Requirement ID	[rte_sws_in_0077]
Object identifier	RunnableEntityRunsInExclusiveArea
Object information	the attribute insideExclusiveArea of Runnable Entity
Description	The RunnableEntity is inside the referenced ExclusiveArea
Rationale	Defines the exclusive area the RunnableEntity is in. It is necessary for consistency mechanisms.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::InternalBehavior::RunnableEntity::insideExclusiveArea
Required by	rte_sws_3500
Contract phase	No

Requirement ID	[rte_sws_in_3017]
Object identifier	IrvCommAppr
Object information	Communication approach of InterRunnableVariable
Description	Whether the access to the InterRunnableVariable is explicit or implicit
Rationale	For generation of the API for accessing the InterRunnableVariable.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::InternalBehavior::InterRunnableVariable::communicationApproach
Required by	rte_sws_3580
Contract phase	Yes

Requirement ID	[rte_sws_in_0078]
Object identifier	ReadVariable
Object information	the attribute readVariable of Runnable Entity
Description	Inter-runnable variables to which this RunnableEntity has implicit read access.
Rationale	It is necessary for consistency mechanisms.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::InternalBehavior::RunnableEntity::readVariable

Required by	rte_sws_1303
Contract phase	Yes

Requirement ID	[rte_sws_in_0082]
Object identifier	WriteVariable
Object information	the attribute writeVariable of Runnable Entity
Description	Inter-runnable variables to which this RunnableEntity has implicit write access.
Rationale	It is necessary for consistency mechanisms.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::InternalBehavior::RunnableEntity::writeVariable
Required by	rte_sws_1304
Contract phase	Yes

Requirement ID	[rte_sws_in_0080]
Object identifier	RunnableEntityCanEnterExclusiveArea
Object information	the attribute usesExclusiveArea of Runnable Entity
Description	The RunnableEntity can enter and exit the referenced exclusive area
Rationale	It is necessary for the data consistency mechanism.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::InternalBehavior::RunnableEntity::usesExclusiveArea
Required by	rte_sws_1307 rte_sws_1308
Contract phase	No

B.6 RTE configuration

Requirement ID	[rte_sws_in_0037]
Object identifier	CompatibilityMode
Object information	RTE generation compatibility mode
Description	RTE generation mode that ensures RTE API compatibility on object code level.
Rationale	The compatibility mode shall be supported by all RTE generators
Template meta-model path	M1::AUTOSAR Descriptions::ECUCParameterDefinition::RTE::GenerationParameters::RteGenerationMode::CompatibilityMode
Required by	rte_sws_1151
Contract phase	Yes

Requirement ID	[rte_sws_in_0038]
Object identifier	VendorMode

Object information	RTE generation vendor mode
Description	RTE generation mode that provides an vendor-specific optimized RTE implementation
Rationale	An RTE generator may optionally support vendor mode. RTE generators from different vendors are unlikely to be compatible when run in the vendor mode
Template meta-model path	M1::AUTOSAR Descriptions::ECUCParameterDefinition::RTE::GenerationParameters::RteGenerationMode::VendorMode
Required by	rte_sws_1152
Contract phase	Yes

Requirement ID	[rte_sws_in_5018]
Object identifier	RteVfbTrace
Object information	Enable VFB tracing
Description	RTE generator will generate code to trace the communication on certain VFB communication
Rationale	The RTE generator shall be able to enable/disable VFB tracing.
Template meta-model path	M1::AUTOSAR Descriptions::ECUCParameterDefinition::RTE::GenerationParameters::RteVfbTrace
Required by	rte_sws_1322 rte_sws_1323 rte_sws_1327 rte_sws_1328
Contract phase	No

Requirement ID	[rte_sws_in_5019]
Object identifier	RteVfbTraceFunction
Object information	VFB tracing hook functions
Description	RTE generator will generate VFB tracing calls only for the defined communications.
Rationale	To be able to select which communication should be traced.
Template meta-model path	M1::AUTOSAR Descriptions::ECUCParameterDefinition::RTE::GenerationParameters::RteVfbTraceFunction
Required by	rte_sws_1324 rte_sws_1325
Contract phase	No

Requirement ID	[rte_sws_in_5060]
Object identifier	RteOptimizationMode
Object information	Rte Generator optimization mode
Description	RTE Generator will optimize for Memory or Runtime

Rationale	During RTE Generation several decisions have to be taken which influence the memory and runtime consumption of the generated RTE. This switch forces the RTE Generator to apply one optimization direction.
Template meta-model path	M1::AUTOSAR Descriptions::ECUCParameterDefinition::RTE::GenerationParameters::RteOptimizationMode
Required by	rte_sws_5053
Contract phase	No

B.7 Measurement and calibration

Requirement ID	[rte_sws_in_3944]
Object identifier	MeasurementSupport
Object information	Measurement support enabling
Description	The RTE generator shall have the option to switch measurement support on and off
Rationale	Measurement is mainly needed for development and when enabled prohibits some RAM usage optimization
Template meta-model path	M1::AUTOSAR Descriptions::ECUCParameterDefinition::RTE::GenerationParameters::RteMeasurementSupport
Required by	rte_sws_3903
Contract phase	No

Requirement ID	[rte_sws_in_3945]
Object identifier	CalibrationSupport
Object information	Calibration support enabling
Description	The RTE generator shall support data emulation without SW support and several methods of data emulation with SW support
Rationale	Data emulation without SW support is used when special calibration support HW is present. Without HW support different project needs require different data emulation with SW support methods
Template meta-model path	M1::AUTOSAR Descriptions::ECUCParameterDefinition::RTE::GenerationParameters::RteCalibrationSupport
Required by	rte_sws_3942 rte_sws_3910 rte_sws_3943
Contract phase	No

Requirement ID	[rte_sws_in_3946]
Object identifier	ComponentCalibrationSupport
Object information	Granularity of calibration support per CalprmComponent instance
Description	Separate calibration support enabling for each CalprmComponentType instance and AUTOSAR SWC

Rationale	Project specific needs exist what to calibrate. Partial enabling saves resources for data emulation with SW support
Template meta-model path	M1::AUTOSAR Descriptions::ECUCParameterDefinition::RTE::ComponentTypeCalibration::CalibrationSupportEnabled
Required by	rte_sws_3905 rte_sws_3906
Contract phase	No

Requirement ID	[rte_sws_in_5048]
Object identifier	swAddrMethod
Object information	CalibrationCategory
Description	The RTE generator shall separate calibration parameters from Calprm-ComponentPrototypes respectively AUTOSAR SW-Cs depending on the CalprmElementPrototype property swAddrMethod.
Rationale	
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::Datatype::Datatypes::DataPrototype::SWDataDefProps::swAddrMethod
Required by	rte_sws_3907
Contract phase	No

B.8 Mode management

Requirement ID	[rte_sws_in_2688]
Object identifier	ModeDeclarationGroup
Description	The ModeDeclarationGroup provides an abstract definition of a mode machine (state machine) of non-overlapping modes. It contains ModeDeclarations that represent the modes.
Rationale	The ModeDeclarationGroup is needed for the type definitions required to represent the modes of a mode machine.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::ModeDeclaration::ModeDeclarationGroup
Required by	rte_sws_2542, rte_sws_2627, rte_sws_2659
Contract phase	YES

Requirement ID	[rte_sws_in_2689]
Object identifier	ModeDeclaration
Description	The ModeDeclaration represents one mode of a ModeDeclarationGroup .
Rationale	The ModeDeclarations are needed to define ModeDisablingDependencies and ModeSwitchEvents . RTE uses the short name of ModeDeclarations for the symbol definition to identify the modes.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::ModeDeclaration::ModeDeclaration

Required by	rte_sws_2542, rte_sws_2567, rte_sws_2546, rte_sws_2558, rte_sws_2631, rte_sws_2660
Contract phase	YES

Requirement ID	[rte_sws_in_2690]
Object identifier	initialMode
Description	The initialMode is a reference of the ModeDeclarationGroup to it's initial mode.
Rationale	The initialMode is needed to define the mode of each <code>rte_sws_mode</code> machine instance after startup of the RTE.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::ModeDeclaration::ModeDeclarationGroup.initialMode
Required by	rte_sws_2544
Contract phase	No

Requirement ID	[rte_sws_in_2691]
Object identifier	ModeDeclarationGroupPrototype
Description	A ModeDeclarationGroupPrototype is used in sender receiver interfaces. All connected ports of compatible interfaces with the same ModeDeclarationGroupPrototype instantiate a <code>mode machine instance</code> .
Rationale	The ModeDeclarationGroupPrototype is needed for the instantiation of mode machines and for the communication using <code>mode ports</code> .
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::PortInterface::ModeDeclarationGroupPrototype
Required by	rte_sws_2630, rte_sws_2549, rte_sws_2546
Contract phase	YES

Requirement ID	[rte_sws_in_2692]
Object identifier	ModeSwitchEvent
Description	The ModeSwitchEvent describes the event of entering or leaving a mode.
Rationale	The ModeSwitchEvent is needed to describe the triggering of a runnable by a mode switch. The ModeSwitchEvent requires the attribute activation to define, if the event is triggered on entering or leaving the mode and the reference startOnEvent to define the runnable, it triggers.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::InternalBehavior::RteEvents::ModeSwitchEvent
Required by	rte_sws_2562, rte_sws_2564
Contract phase	No

Requirement ID	[rte_sws_in_0036]
----------------	--------------------------

Object identifier	ModeDisablingDependency
Object information	Dependency between modes and disabling of RTEEvents .
Description	The ModeDisablingDependency describes the mode disabling of an RTEEvent as a trigger for a runnable. It belongs to an RTEEvent and requires a reference dependentOnMode to the instance of a ModeDeclaration which shall result in the disabling.
Rationale	The existence of a ModeDisablingDependency shall prevent the RTE to start a runnable by the corresponding event in the referenced mode
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::ModeDeclaration::ModeDisablingDependency
Required by	rte_sws_2503, rte_sws_2661, rte_sws_2663
Contract phase	No

Requirement ID	[rte_sws_in_2693]
Object identifier	ModeSwitchComSpec.queueLength
Object information	The ModeSwitchComSpec.queueLength is an attribute of the ModeSwitchComSpec of a provide mode port.
Description	The ModeSwitchComSpec.queueLength defines the size of the input queue of mode switch notifications to a mode machine. Only the ComSpec of an AtomicSoftwareComponentType PortPrototype shall be considered.
Rationale	Needed to configure RTE's queues for mode switches.
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::Communication::ModeSwitchComSpec.queueLength
Required by	rte_sws_2667, rte_sws_2668, rte_sws_2624, rte_sws_2675, rte_sws_2672
Contract phase	No

Requirement ID	[rte_sws_in_2694]
Object identifier	ModeSwitchedAckRequest
Object information	The ModeSwitchedAckRequest is an element of the are attributes of the ModeSwitchComSpec of a provide mode port.
Description	The ModeSwitchedAckRequest indicates that the mode manager needs a feedback of the completion of the notified mode switches. ModeSwitchedAckRequest contains an attribute timeout for the feedback. timeout = 0 configures no timeout.
Rationale	Needed for the configuration of the Rte_Feedback API and of the ModeSwitchedAckEvent
Template meta-model path	M2::AUTOSAR Templates::SWComponentTemplate::Communication::ModeSwitchedAckRequest
Required by	rte_sws_2587
Contract phase	Yes

Requirement ID	[rte_sws_in_2695]
Object identifier	ModeSwitchedAckEvent
Description	The ModeSwitchedAckEvent represents an event that is available to the mode manager after the completion of a mode switch. It contains a reference 'eventSource' to reference the ModeSwitchPoint to which the acknowledgement belongs. It can be used with a WaitPoint to generate a blocking <code>Rte_Feedback</code> API or with a startOnEvent reference to trigger a runnable.
Rationale	Needed to trigger the reception of a mode switch completion acknowledgement by the mode manager.
Template meta-model path	M2::AUTOSAR_Templates::SWComponentTemplate::InternalBehavior::RteEvents::ModeSwitchedAckEvent
Required by	rte_sws_2587
Contract phase	No

Requirement ID	[rte_sws_in_2696]
Object identifier	ModeSwitchPoint
Description	A ModeSwitchPoint represents the position within the <code>mode manager</code> where the mode switch is initiated.
Rationale	The ModeSwitchPoint is required to define the runnable, that may use the <code>Rte_Switch</code> API.
Template meta-model path	M2::AUTOSAR_Templates::SWComponentTemplate::InternalBehavior::ModeDeclarationGroup::ModeSwitchPoint
Required by	rte_sws_
Contract phase	YES/No

C External Requirements

[rte_sws_ext_2054] The RTE-Generator expects only one instance of the ECU Abstraction.

[rte_sws_ext_7001] The runnables worst case execution time shall be less than the GCD of all runnables period and offset in activation offset context for runnables.

[rte_sws_ext_2559] The RTE configurator shall have access to the schedule table configuration (see also `rte_sws_4014`)

[rte_sws_ext_2542] Whenever any runnable entity is running, there shall always be exactly one mode or one mode transition active of each ModeDeclarationGroup.

[rte_sws_ext_2507] The mode switch shall be notified to the mode user (and RTE) locally on each ECU.

[rte_sws_ext_3813] The indirect API may only be used for a port if explicit data access to this port is specified via `DataSendPoints` and `DataReceivePoints`.

[rte_sws_ext_2680] The `Rte_Send/Rte_Write` APIs may only be used by the runnable that contains the corresponding `DataSendPoint`

[rte_sws_ext_2681] The `Rte_Switch` API may only be used by the runnable that contains the corresponding `ModeSwitchPoint`

[rte_sws_ext_2682] The `Rte_Invalidate` API may only be used by the runnable that contains the corresponding `DataSendPoint`

[rte_sws_ext_2687] A blocking `Rte_Feedback` API may only be used by the runnable that contains the corresponding `WaitPoint`

[rte_sws_ext_2683] The `Rte_Read` API may only be used by the runnable that contains the corresponding `DataReceivePoint`

[rte_sws_ext_2684] The `Rte_Receive` API may only be used by the runnable that contains the corresponding `DataReceivePoint`

[rte_sws_ext_2685] The `Rte_Call` API may only be used by the runnable that contains the corresponding `ServerCallPoint`

[rte_sws_ext_2686] The blocking `Rte_Result` API may only be used by the runnable that contains the corresponding `WaitPoint`

[rte_sws_ext_2601] The `Rte_IStatus` API shall only be used by a `RunnableEntity` that either has a `DataReadAccess` referring to the `DataElementPrototype` or is

triggered by a `DataReceiveErrorEvent` referring to the `DataElementPrototype`.

[rte_sws_ext_2704] Only the least significant six bit of the return value of a server runnable shall be used by the application to indicate an error. The upper two bit shall be zero.

[rte_sws_ext_2582] `Rte_Start` shall be called only once by the `EcuStateManager` after the basic software modules required by RTE are initialized.

[rte_sws_ext_2583] `Rte_Stop` shall be called by the `EcuStateManager` before the basic software modules required by RTE are shut down.

D MISRA C Compliance

In general, all RTE code, whether generated or not, shall conform to the HIS subset of the MISRA C standard `rte_sws_1168` [18]. This chapter lists all the MISRA C rules of the HIS subset that may be violated by the generated RTE.

The MISRA C standard was defined with having mainly hand-written code in mind. Part of the MISRA C rules only apply to hand-written code, they do not make much sense in the context of automatic code generation. Additionally, there are some rules that are violated because of technical reasons, mainly to reduce RTE overhead.

The rules listed in this chapter are expected to be violated by RTE code. Violations to the rules listed here do not need to be documented as non-compliant to MISRA C in the generated code itself.

MISRA rule	11
Description	Identifiers (internal and external) shall not rely on significance of more than 31 characters. Furthermore the compiler/linker shall be checked to ensure that 31 character significance and case sensitivity are supported for external identifiers.
Violations	The defined RTE naming convention may result in identifiers with more than 31 characters. The compliance to this rule is under user's control.

MISRA rule	23
Description	All declarations at file scope should be static where possible.
Violations	E.g. for the purpose of monitoring during calibration or debugging it may be necessary to use non-static declarations at file scope.

MISRA rule	42
Description	The comma operator shall not be used, except in the control expression of a <i>for</i> loop.
Violations	Function-like macros may have to use the comma operator. Function-like macros are required for efficiency reasons [BSW00330].

MISRA rule	45
Description	Type casting from any type to or from pointers shall not be used.
Violations	For the implementation of exclusive areas (<code>rte_sws_3740</code> , Section 5.4.2.4) casting between pointer types is needed.

MISRA rule	54
Description	A null statement shall only occur on a line by itself, and shall not have any other text on the same line.
Violations	In an optimized RTE, API calls may result in a null statement. Therefore the compliance to this rule cannot be guaranteed.

E Interfaces of COM used by the RTE

The specification of the RTE requires the usage of the following COM API functions and COM callback functions.

COM API function	Context
Com_SendSignal	to transmit a data element of primitive type using COM.
Com_ReceiveSignal	to retrieve the new value of a data element of primitive type from COM.
Com_UpdateShadowSignal	to update a primitive element of a data element of complex type in preparation for sending the complex type using COM.
Com_SendSignalGroup	to initiate sending of a data element of complex type using COM.
Com_ReceiveSignalGroup	to retrieve the new value of a data element of complex type from COM.
Com_ReceiveShadowSignal	to retrieve the new value of a primitive element of a data element of complex type from COM.
Com_InvalidateSignal	to invalidate a data element of primitive type using COM.
Com_InvalidateSignalGroup	to invalidate a whole signal group using COM.

Table E.1: COM API functions used by the RTE

Callback function	Configuration	Usage
Rte_COMCbK_<sn>	COM_NOTIFICATION_SIGNAL of COM_SIGNAL	Notification of data reception of a data element of primitive type
Rte_COMCbKInv_<sn>	COM_RX_DATA_INVALID_INDICATION_FUNCTION of COM_RX_DATA_INVALID of COM_SIGNAL	Notification of reception of an invalidated signal
Rte_COMCbKInv_<sg>	COM_RX_DATA_INVALID_INDICATION_FUNCTION of COM_RX_DATA_INVALID of COM_SIGNAL_GROUP	Notification of reception of an invalidated signal group
Rte_COMCbKTOut_<sn>	COM_NOTIFICATION_ERROR of COM_SIGNAL	Notification of a deadline monitoring violation for a data element of primitive type (only present if aliveTimeout is present)
Rte_COMCbK_<sg>	COM_NOTIFICATION_SIGNAL of COM_SIGNAL_GROUP	Notification of data reception of a data element of complex type

Callback function	Configuration	Usage
Rte_COMCbktOut_<sg>	COM_NOTIFICATION_ERROR of COM_SIGNAL_GROUP	Notification of a deadline monitoring violation for a data element of complex type (only present if aliveTimeout is present)

Table E.2: COM Callback functions provided by the RTE for signal reception

Callback function	Configuration	Usage
Rte_COMCbktAck_<sn>	COM_NOTIFICATION_SIGNAL of COM_SIGNAL	Notification of successful transmission of a data element of primitive type (only present if acknowledgement request is specified)
Rte_COMCbktErr_<sn>	COM_NOTIFICATION_ERROR of COM_SIGNAL	Notification of a transmission error of a data element of primitive type (only present if acknowledgement request is specified)
Rte_COMCbktAck_<sg>	COM_NOTIFICATION_SIGNAL of COM_SIGNAL_GROUP	Notification of successful transmission of a data element of complex type (only present if acknowledgement request is specified)
Rte_COMCbktErr_<sg>	COM_NOTIFICATION_ERROR of COM_SIGNAL_GROUP	Notification of a transmission error of a data element of complex type (only present if acknowledgement request is specified)

Table E.3: COM Callback functions provided by the RTE for signal transmission