

Document Title	Specification of Operating System
Document Owner	AUTOSAR
Document Responsibility	AUTOSAR
Document Identification No	034
Document Classification	Standard

Document Version	5.0.0
Document Status	Final
Part of Release	4.0
Revision	3

Document Change History			
Date	Version	Changed by	Change Description
23.11.2011	5.0.0	AUTOSAR Administration	<ul style="list-style-type: none"> Included MultiCore support from former "Specification of Multi-Core OS Architecture"
22.10.2010	4.1.0	AUTOSAR Administration	<ul style="list-style-type: none"> Clarification in 7.8.1 (meaning of "do nothing") and 7.1.2.1 ("OSEK declarations") Minor changes as typos and rewording
30.11.2009	4.0.0	AUTOSAR Administration	<ul style="list-style-type: none"> Extension of services (Chapter 12) States in OS- Applications Active termination of other OS- Applications in possible (Chapter8) Legal disclaimer revised Chapter 10.4 revised
15.01.2009	3.1.0	AUTOSAR Administration	<ul style="list-style-type: none"> Changes in OS configuration: <ul style="list-style-type: none"> removed "OsAppModeld" Parameter from OsAppModeContainer added optional references from OsAppModeContainer to OsAlarm, OsTask and OsScheduleTable
04.08.2008	3.0.2	AUTOSAR Administration	Legal Disclaimer revised
17.04.2008	3.0.1	AUTOSAR Administration	<ul style="list-style-type: none"> Added "OsScheduleTableDuration" parameter to configuration specification chapter

Document Change History			
Date	Version	Changed by	Change Description
07.12.2007	3.0.0	AUTOSAR Administration	<ul style="list-style-type: none"> • Changed methods for timing protection • Moved configuration from OIL to AUTOSAR XML • Clarified description for synchronization and schedule tables • Document meta information extended • Small layout adaptations made
31.01.2007	2.1.0	AUTOSAR Administration	<ul style="list-style-type: none"> • Added support for SoftwareFreeRunningTimer (SWFRT) incl. 2 new APIs • Added API to start a schedule table synchron • Misc. Corrections, Clarification and further explanations • Legal disclaimer revised • Release Notes added • “Advice for users” revised • “Revision Information” added
28.04.2006	2.0.0	AUTOSAR Administration	<p>Document structure adapted to common Release 2.0 SWS Template.</p> <ul style="list-style-type: none"> • Major changes in chapter 10 • Structure of document changed partly • Other changes see chapter 14
28.06.2005	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 specifications 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 specifications for illustration purposes only, and they themselves are not part of the AUTOSAR Standard. Neither their presence in such specifications, 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 Content

1	Introduction and functional overview	11
2	Acronyms and abbreviations	12
2.1	Glossary of Terms	12
3	Related documentation.....	17
3.1	Input documents.....	17
3.2	Related standards and norms	18
3.2.1	OSEK/VDX.....	18
3.2.2	HIS.....	18
3.2.3	ISO/IEC.....	19
3.3	Company Reports, Academic Work, etc.....	19
4	Constraints and assumptions	20
4.1	Existing Standards	20
4.2	Terminology	20
4.3	Interaction with the RTE	20
4.4	Operating System Abstraction Layer (OSAL).....	21
4.5	Multi-Core Hardware assumptions	21
4.5.1	CPU Core features.....	21
4.5.2	Memory features	22
4.5.3	Multi-Core Limitations	22
4.6	Limitations	23
4.6.1	Hardware	23
4.6.2	Programming Language.....	23
4.6.3	Miscellaneous	24
4.7	Applicability to car domains.....	24
5	Dependencies to other modules.....	25
5.1	File structure	25
5.1.1	Code file structure.....	25
5.1.2	Header file structure.....	25
6	Requirements Traceability.....	27
7	Functional specification	37
7.1	Core OS	37
7.1.1	Background & Rationale	37
7.1.2	Requirements.....	37
7.1.2.1	Restrictions on OSEK OS	37
7.1.2.2	Undefined Behaviour in OSEK OS.....	38
7.1.2.3	Extensions to OSEK OS	39
7.2	Software Free Running Timer	40
7.3	Schedule Tables.....	41
7.3.1	Background & Rationale	41
7.3.2	Requirements.....	41
7.3.2.1	Structure of a Schedule Table.....	41

7.3.2.2	Constraints on Expiry Points	42
7.3.2.3	Processing Schedule Tables.....	43
7.3.2.4	Repeated Schedule Table Processing.....	44
7.3.2.5	Controlling Schedule Table Processing	45
7.4	Schedule Table Synchronization	48
7.4.1	Background & Rationale	48
7.4.2	Requirements.....	50
7.4.2.1	Implicit Synchronization	50
7.4.2.2	Explicit Synchronization	51
7.4.2.3	Performing Synchronization.....	55
7.5	Stack Monitoring Facilities.....	57
7.5.1	Background & Rationale	57
7.5.2	Requirements.....	58
7.6	OS-Application	58
7.6.1	Background & Rationale	58
7.6.2	Requirements.....	60
7.7	Protection Facilities	62
7.7.1	Memory Protection	62
7.7.1.1	Background & Rationale	62
7.7.1.2	Requirements.....	63
7.7.2	Timing Protection	65
7.7.2.1	Background & Rationale	65
7.7.2.2	Requirements.....	68
7.7.2.3	Implementation Notes	70
7.7.3	Service Protection	71
7.7.3.1	Invalid Object Parameter or Out of Range Value	71
7.7.3.2	Service Calls Made from Wrong Context	72
7.7.3.3	Services with Undefined Behaviour	73
7.7.3.4	Service Restrictions for Non-Trusted OS-Applications.....	75
7.7.3.5	Service Calls on Objects in Different OS-Applications	76
7.7.4	Protecting the Hardware used by the OS.....	77
7.7.4.1	Background & Rationale	77
7.7.4.2	Requirements.....	77
7.7.4.3	Implementation Notes	77
7.7.5	Providing »Trusted Functions«.....	78
7.7.5.1	Background & Rationale	78
7.7.5.2	Requirements.....	78
7.8	Protection Error Handling	78
7.8.1	Background & Rationale	79
7.8.2	Requirements.....	80
7.9	Operating System for Multi-Core.....	81
7.9.1	Background & Rationale	82
7.9.1.1	Requirements.....	82
7.9.2	Scheduling	82
7.9.2.1	Requirements.....	83
7.9.3	Locatable entities (LE)	83
7.9.3.1	Requirements.....	83
7.9.4	Multi-Core start-up concept.....	84
7.9.4.1	Requirements.....	86

7.9.5	Cores under control of the AUTOSAR OS	87
7.9.5.1	Requirements.....	87
7.9.6	Cores which are not controlled by the AUTOSAR OS.....	88
7.9.6.1	Requirements.....	88
7.9.7	Multi-Core shutdown concept.....	88
7.9.7.1	Synchronized shutdown concept	88
7.9.7.2	Individual shutdown concept.....	89
7.9.7.3	Shutdown in case of fatal internal errors.....	89
7.9.8	OS service functionality (overview)	90
7.9.9	GetTaskID.....	91
7.9.10	Interrupt disabling.....	92
7.9.10.1	Requirements.....	92
7.9.11	TASK activation.....	92
7.9.11.1	Requirements.....	92
7.9.12	TASK Chaining.....	93
7.9.12.1	Requirements.....	93
7.9.13	EVENT setting.....	93
7.9.13.1	Requirements.....	93
7.9.14	Activating additional cores	94
7.9.15	Start of the OS	94
7.9.15.1	Requirements.....	94
7.9.16	TASK termination	95
7.9.16.1	Requirements.....	95
7.9.17	Termination of OS-Applications.....	95
7.9.17.1	Requirements.....	95
7.9.18	Shutdown of the OS	96
7.9.18.1	Requirements.....	96
7.9.19	Waiting for EVENTS	96
7.9.19.1	Requirements.....	97
7.9.20	Calling trusted functions	97
7.9.20.1	Requirements.....	97
7.9.21	Invoking reschedule	97
7.9.21.1	Requirements.....	97
7.9.22	RESOURCE occupation	97
7.9.23	The CoreID.....	98
7.9.23.1	Requirements.....	98
7.9.24	COUNTERs, background & rationale.....	99
7.9.25	Multi-Core restrictions on COUNTERs.....	99
7.9.25.1	Requirements.....	99
7.9.26	Synchronization of COUNTERs	100
7.9.27	ALARMS.....	101
7.9.27.1	Requirements.....	101
7.9.28	Schedule tables.....	102
7.9.28.1	Requirements.....	102
7.9.29	The spinlock mechanism.....	102
7.9.29.1	Requirements.....	104
7.9.30	Offline checks.....	105
7.9.30.1	Requirements.....	106
7.9.31	Auto start Objects.....	106

7.9.31.1	Requirements.....	106
7.10	Inter-OS-Application Communicator (IOC)	107
7.10.1	Background & Rationale	107
7.10.2	IOC - General purpose.....	108
7.10.3	IOC functionality	109
7.10.3.1	Communication	109
7.10.3.2	Notification	109
7.10.4	IOC interface.....	110
7.10.5	IOC internal structure	111
7.10.6	IOC configuration and generation	111
7.10.7	IOC integration examples.....	112
7.10.7.1	Example 1 - 1:1 sender/receiver communication without notification 112	
7.10.7.2	Example 2 - N:1 client/server communication with receiver notification by RTE.....	114
7.10.8	Future extensions.....	115
7.11	System Scalability	115
7.11.1	Background & Rationale	115
7.11.2	Requirements.....	117
7.12	Hook Functions	118
7.12.1	Background & Rationale	118
7.12.2	Requirements.....	118
7.13	Error classification.....	119
7.14	Debug support.....	120
8	API specification.....	121
8.1	Constants	121
8.1.1	Error codes of type StatusType.....	121
8.2	Macros	121
8.3	Type definitions	121
8.3.1	ApplicationType (for OS-Applications)	121
8.3.2	ApplicationStateType	122
8.3.3	ApplicationStateRefType.....	122
8.3.4	TrustedFunctionIndexType	122
8.3.5	TrustedFunctionParameterRefType	122
8.3.6	AccessType.....	122
8.3.7	ObjectAccessType	122
8.3.8	ObjectTypeType.....	122
8.3.9	MemoryStartAddressType.....	123
8.3.10	MemorySizeType	123
8.3.11	ISRType	123
8.3.12	ScheduleTableType	123
8.3.13	ScheduleTableStatusType	124
8.3.14	ScheduleTableStatusRefType.....	124
8.3.15	CounterType	124
8.3.16	ProtectionReturnType	124
8.3.17	RestartType.....	124
8.3.18	PhysicalTimeType.....	125
8.3.19	CoreIDType.....	125

8.3.20	SpinlockIdType.....	125
8.3.21	TryToGetSpinlockType	125
8.4	Function definitions	125
8.4.1	GetApplicationID	125
8.4.2	GetISRID.....	126
8.4.3	CallTrustedFunction	127
8.4.4	CheckISRMemoryAccess	129
8.4.5	CheckTaskMemoryAccess.....	129
8.4.6	CheckObjectAccess	130
8.4.7	CheckObjectOwnership	131
8.4.8	StartScheduleTableRel	132
8.4.9	StartScheduleTableAbs	133
8.4.10	StopScheduleTable	134
8.4.11	NextScheduleTable	135
8.4.12	StartScheduleTableSynchron.....	137
8.4.13	SyncScheduleTable	138
8.4.14	SetScheduleTableAsync	139
8.4.15	GetScheduleTableStatus	140
8.4.16	IncrementCounter	141
8.4.17	GetCounterValue	142
8.4.18	GetElapsedValue	143
8.4.19	TerminateApplication	144
8.4.20	AllowAccess	145
8.4.21	GetApplicationState	146
8.4.22	GetNumberOfActivatedCores.....	146
8.4.23	GetCoreID	147
8.4.24	StartCore.....	147
8.4.25	StartNonAutosarCore.....	148
8.4.26	GetSpinlock.....	149
8.4.27	ReleaseSpinlock	150
8.4.28	TryToGetSpinlock	151
8.4.29	ShutdownAllCores.....	153
8.5	IOC.....	153
8.5.1	Imported types	153
8.5.2	Type definitions	154
8.5.3	Constants	154
8.5.4	Function definitions	155
8.5.4.1	locSend/locWrite.....	155
8.5.4.2	locSendGroup/locWriteGroup.....	157
8.5.4.3	locReceive/locRead.....	160
8.5.4.4	locReceiveGroup/locReadGroup	162
8.5.4.5	locEmptyQueue	164
8.6	Expected Interfaces.....	164
8.6.1	Mandatory Interfaces	164
8.6.2	Optional Interfaces	165
8.6.2.1	ReceiverPullICB.....	165
8.7	Hook functions.....	166
8.7.1	Protection Hook.....	166
8.7.2	Application specific StartupHook.....	166

8.7.3	Application specific ErrorHook	167
8.7.4	Application specific ShutdownHook	167
9	Sequence diagrams.....	169
9.1	Sequence chart for calling trusted functions.....	169
9.2	Sequence chart for usage of ErrorHook	170
9.3	Sequence chart for ProtectionHook.....	171
9.4	Sequence chart for StartupHook	172
9.5	Sequence chart for ShutdownHook.....	173
9.6	Sequence diagrams of Sender Receiver communication over the IOC....	173
9.6.1	LastIsBest communication	173
9.6.2	Queued communication without pull callback.....	174
9.6.3	Queued communication with pull callback	176
10	Configuration Specification	177
10.1	How to read this chapter	177
10.1.1	Configuration and configuration parameters	177
10.1.2	Variants.....	177
10.1.3	Containers.....	178
10.1.4	Rules for paramters.....	178
10.2	Containers and configuration parameters	178
10.2.1	Variants.....	178
10.2.2	Os	178
10.2.3	OsAlarmSetEvent.....	179
10.2.4	OsAlarm	180
10.2.5	OsAlarmAction	181
10.2.6	OsAlarmActivateTask.....	181
10.2.7	OsAlarmAutostart.....	181
10.2.8	OsAlarmCallback	182
10.2.9	OsAlarmIncrementCounter	183
10.2.10	OsApplication	183
10.2.11	OsApplicationHooks	185
10.2.12	OsApplicationTrustedFunction.....	186
10.2.13	OsAppMode.....	187
10.2.14	OsCounter	187
10.2.15	OsEvent.....	189
10.2.16	OsHooks.....	190
10.2.17	Oslsr	191
10.2.18	OslsrResourceLock.....	192
10.2.19	OslsrTimingProtection	192
10.2.20	OsOS.....	194
10.2.21	OsResource.....	196
10.2.22	OsScheduleTable	197
10.2.23	OsScheduleTableAutostart.....	198
10.2.24	OsScheduleTableEventSetting.....	199
10.2.25	OsScheduleTableExpiryPoint	199
10.2.26	OsScheduleTableTaskActivation	200
10.2.27	OsScheduleTblAdjustableExpPoint	200
10.2.28	OsScheduleTableSync	201
10.2.29	OsSpinlock	202

10.2.30	OsTask	203
10.2.31	OsTaskAutostart.....	204
10.2.32	OsTaskResourceLock	205
10.2.33	OsTaskTimingProtection	205
10.2.34	OsTimeConstant.....	207
10.3	Containers and configuration parameter extensions of the IOC.....	207
10.3.1	Osloc.....	207
10.3.2	OslocCommunication.....	208
10.3.3	OslocSenderProperties.....	209
10.3.4	OslocReceiverProperties	210
10.3.5	OslocDataProperties.....	211
10.4	Published Information.....	212
11	Generation of the OS	213
11.1	Read in configuration	213
11.2	Consistency check	213
11.3	Generating operating system	215
12	Application Notes.....	216
12.1	Hooks.....	216
12.2	Providing Trusted Functions.....	216
12.3	Migration hints for OSEKtime OS users	218
12.4	Software Components and OS-Applications	220
12.5	Global Time Synchronization	221
12.6	Working with FlexRay.....	221
12.7	Migration from OIL to XML	222
12.8	Migrating RES_SCHEDULER in AUTOSAR OS.....	222
12.9	Debug support.....	223
12.10	Integration hints for peripheral protection	223
13	AUTOSAR Service implemented by the OS	225
13.1	Scope of this Chapter.....	225
13.1.1	Package.....	225
13.2	Overview	225
13.3	Specification of the Ports and Port Interfaces	226
13.3.1	Data Types and Port Interface	226
13.3.1.1	General Approach.....	226
13.3.1.2	Data Types.....	226
13.3.1.3	Port Interface	226
13.3.1.4	Ports	227
14	Outlook on Memory Protection Configuration	228
14.1	Configuration Approach.....	228
15	Changes to Release 3.0/3.1	229
16	Not applicable requirements	230

1 Introduction and functional overview

This document describes the essential requirements on the AUTOSAR Operating System to satisfy the top-level requirements presented in the AUTOSAR SRS [2].

In general, operating systems can be split up in different groups according to their characteristics, e.g. statically configured vs. dynamically managed. To classify the AUTOSAR OS, here are the basic features: the OS

- is configured and scaled statically
- is amenable to reasoning of real-time performance
- provides a priority-based scheduling policy
- provides protective functions (memory, timing etc.) at run-time
- is hostable on low-end controllers and without external resources

This feature set defines the type of OS commonly used in the current generation of automotive ECUs, with the exception of Telematic/Infotainment systems. It is assumed that Telematic/Infotainment systems will continue to use proprietary Oss under the AUTOSAR framework (e.g. Windows CE, VxWorks, QNX, etc.). In the case where AUTOSAR components are needed to run on these proprietary Oss, the interfaces defined in this document should be provided as an Operating System Abstraction Layer (OSAL).

This document uses the industry standard OSEK OS [15] (ISO 17356-3) as the basis for the AUTOSAR OS. The reader should be familiar with this standard before reading this document.

This document describes extensions to, and restrictions of, this OSEK OS.

2 Acronyms and abbreviations

Abbreviation	Description
API	Application Programming Interface
AR	AUTOSAR
BSW	Basic Software
BSWMD	Basic Software Module Description
CDD	Complex Device Driver
COM	Communication
ECC	Extended Conformance Class
ECU	Electronic Control Unit
HIS	Hersteller Initiative Software
HW	Hardware
ID	Identifier
IOC	Inter OS-Application communicator
ISR	Interrupt Service Routine
LE	A locatable entity is a distinct piece of software that has the same effect regardless of which core it is located.
MC	Multi-Core
MCU	Microcontroller Unit
ME	Mutual exclusion
MPU	Memory Protection Unit
NMI	Mutual exclusion
OIL	OSEK Implementation Language
OS	Operating System
OSEK/VDX	Offene Systeme und deren Schnittstellen für die Elektronik im Kraftfahrzeug
RTE	Run-Time Environment
RTOS	Real Time Operating System
SC	Single-Core
SLA	Software Layered Architecture
SW	Software
SWC	Software Component
SWFRT	Software FreeRunningTimer

2.1 Glossary of Terms

Term:	Definition	
Access Right	An indication that an object (e.g. Task, ISR, hook function) of an OS-Application has the permission of access or manipulation with respect to memory, OS services or (set of) OS objects.	
Cardinality	The number of items in a set.	
Counter	An operating system object that registers a count in ticks. There are two types of counters:	
	Hardware Counter	A counter that is advanced by hardware (e.g. timer). The count value is maintained by the peripheral "in hardware".
	Software Counter	A counter which is incremented by making the <code>IncrementCounter()</code> API call (see OS399). The count value is maintained by the operating system "in software".
Deadline	The time at which a Task/Category 2 ISR must reach a certain point during its execution defined by system design relative to the stimulus that triggered activation. See Figure 2.1	

Delay	<p>The number of ticks between two adjacent expiry points on a schedule table. A pair of expiry points X and Y are said to be adjacent when:</p> <ul style="list-style-type: none"> There is no expiry point Z such that $X.Offset < Z.Offset < Y.Offset$. In this case the $Delay = Y.Offset - X.Offset$ X and Y are the Final Expiry Point and the Initial Expiry Point respectively. In this case $Delay = (Duration - X.Offset) + Y.Offset$ <p>When used in the text, Delay is a relative number of ticks measured from a specified expiry point. For example: X.Delay is the delay from X to the next expiry point.</p>	
Deviation	<p>The minimum number of ticks between the current position on an explicitly synchronized schedule table and the value of the synchronization count modulo the duration of the schedule table.</p>	
Duration	<p>The number of ticks from a notional zero at which a schedule table wraps.</p>	
Execution Time	<p>Tasks:</p> <p>The net time a task spends in the <code>RUNNING</code> state without entering the <code>SUSPENDED</code> or <code>WAITING</code> state excluding all preemptions due to ISRs which preempt the task. An extended task executing the <code>waitEvent()</code> API call to wait on an event which is already set notionally enters the <code>WAITING</code> state. For multiple activated basic tasks the net time is per activation of a task.</p> <p>ISRs:</p> <p>The net time from the first to the last instruction of the user provided Category 2 interrupt handler excluding all preemptions due to higher priority ISRs executing in preference.</p> <p>Execution time includes the time spent in the error, pretask and posttask hooks and the time spent making OS service calls.</p>	
Execution Budget	<p>Maximum permitted execution time for a Task/ISR.</p>	
Expiry Point	<p>The offset on a Schedule Table, measured from zero, at which the OS activates tasks and/or sets events.</p>	
	Initial Expiry Point	The expiry point with the smallest offset
	Final Expiry Point	The expiry point with the largest offset
Hook Function	<p>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</p>	
	Application-specific	Hook functions within the scope of an individual OS-Application.
	System-specific	Hook functions within the scope of the complete system (in general provided by the integrator).
Initial Offset	<p>The smallest expiry point offset on a schedule table. This can be zero.</p>	
Interarrival Time	<p>Basic Tasks</p> <p>The time between successively entering the <code>READY</code> state from the <code>SUSPENDED</code> state. Activation of a task always represents a new arrival. This applies in the case of multiple activations, even if an existing instance of the task is in the <code>RUNNING</code> or <code>READY</code> state.</p> <p>Extended Tasks:</p> <p>The time between successively entering the <code>READY</code> state from the <code>SUSPENDED</code> or <code>WAITING</code> states. Setting an event for a task in the <code>WAITING</code> state represents a new arrival if the task is waiting on the event. Waiting for an event in the <code>RUNNING</code> state which is already set represents a new arrival.</p> <p>ISRs:</p> <p>The time between successive occurrences of an interrupt.</p> <p>See Figure 2.1.</p>	

Interrupt Lock Time	The time for which a Task/ISR executes with Category 1 interrupts disabled/suspended and/or Category 2 interrupts disabled/suspended .	
Interrupt Source Enable	The switch which enables a specific interrupt source in the hardware.	
Interrupt Vector Table	Conceptually, the interrupt vector table contains the mapping from hardware interrupt requests to (software) interrupt service routines. The real content of the Interrupt Vector Table is very hardware specific, e.g. it can contain the start addresses of the interrupt service routines.	
Final Delay	The difference between the Final Expiry Point offset and the duration on a schedule table in ticks. This value defines the delay from the Final Expiry Point to the logical end of the schedule table for single-shot and "nexted" schedule tables.	
Forced OS-Application Termination	The operating system frees all system objects, e.g. forcibly terminates Tasks, disables interrupts, etc., which are associated to the OS-Application. OS-Application and internal variables are potentially left in an undefined state.	
Forced Termination	The OS terminates the Task/Category 2 ISR and does "unlock" its held resources. For details see OS108 and OS109 .	
Linker File	File containing linking settings for the linker. The syntax of the linker file depends on the specific linker and, consequently, definitions are stored "linker-specific" in the linker file.	
Lock Budget	Maximum permitted Interrupt Lock Time or Resource Lock Time.	
Master core	A master core is a core from which the AUTOSAR system is bootstrapped.	
Memory Protection Unit	A Memory Protection Unit (MPU) enables memory partitioning with individual protection attributes. This is distinct from a Memory Management Unit (MMU) that provides a mapping between virtual addresses and physical memory locations at runtime. Note that some devices may realise the functionality of an MPU in an MMU.	
Mode	Describes the permissions available on a processor.	
	Privileged	In general, in »privileged mode« unrestricted access is available to memory as well as the underlying hardware.
	Non-privileged	In »non-privileged mode« access is restricted.
Modulus	The number of ticks required to complete a full wrap of an OSEK counter. This is equal to <code>OsCounterMaxAllowedValue + 1</code> ticks of the counter.	
OS-Application	A collection of OS objects	
	Trusted	An OS-Application that is executed in privileged mode and has unrestricted access to the API and hardware resources. Only trusted applications can provide trusted functions.
	Non-trusted	An OS-Application that is executed in non-privileged mode has restricted access to the API and hardware resources.
OS object	Object that belongs to a single OS-Application: Task, ISR, Alarm, Event, Schedule Table, Resource, Trusted Function, Counter, Applicaton-specific hook.	
OS Service	OS services are the API of the operating system.	
Protection Error	Systematic error in the software of an OS-Application.	
	Memory access violation	A protection error caused by access to an address in a manner for which no access right exists.
	Timing fault	A protection error that violates the timing protection.
	Illegal service	A protection error that violates the service protection, e.g. unauthorized call to OS service.
	Hardware exception	division by zero, illegal instruction etc.
Resource Lock Time	The time an OSEK resource is held by a Task/ISR (excluding the preemptions of the Task/ISR by higher prior Tasks/ISRs).	
Response Time	The time between a Task/ISR being made ready to execute and generating a specified response. The time includes all preemptions. See Figure 2.1	
Restart an OS-Application	An OS-Application can be restarted after self-termination or being forcibly terminated because of a protection error. When an OS-Application is restarted, the OS activates the configured <code>OsRestartTask</code> .	
Scalability Class	The features of the OS (e.g. Memory Protection or Timing Protection), described	

	by this document, can be grouped together to customize the operating system to the needs of the application. There are 4 defined groups of features which are named scalability classes. For details see Chapter 7.11
Schedule Table	Encapsulation of a statically defined set of expiry points.
Section	Part of an object file in which instructions or data are combined to form a unit (contiguous address space in memory allocated for data or code). A section in an object file (object file format) has a name and a size. From the linker perspective, two different sides can be distinguished:
	Input section memory section in an input object file of the linker.
	Output section memory section in an output object file of the linker.
Set (of OS objects)	This document uses the term set, indicating a collection of the same type of OS objects, in the strict mathematical sense, i.e.: - a set contains zero or more OS objects (this means a set can be empty) - the OS objects in the set are unique (this means there cannot be duplicate OS objects in the set)
Spinlock	A spinlock is a locking mechanism where the TASK waits in a loop ("spins") repeatedly checking for a shared variable to become a certain value. The value indicates whether the lock is free or not. In Multi-Core systems the comparison and changing of the variable typically requires an atomic operation. As the TASK remains active but is not doing anything useful, a spinlock is a busy waiting mechanism
Spinlock variable	A spinlock variable is a shared variable used by a spinlock to indicate whether a spinlock is free or occupied.
Symbol	Address label that can be imported/used by software modules and resolved by the linker. The precise syntax of the labels is linker-specific. Here, these address labels are used to identify the start and end of memory sections.
	Start symbol Tags the start of a memory section
	End symbol Tags the end of a memory section
Synchronization of schedule tables with a synchronization counter	Synchronization with a synchronization counter is achieved, if the expiry points of the schedule table are processed within an absolute deviation from the synchronization counter that is smaller than or equal to a precision threshold.
Synchronization Counter	The "Synchronization Counter", distinct from an OS counter object, is an external counter, external to the OS, against which expiry points of a schedule table are synchronized
Task	A Task is the object which executes (user) code and which is managed by the OS. E.g. the OS switches between different Tasks ("schedules"). There are 2 types of Tasks; for more details see [15].
	Basic Task A Task which can not block by itself. This means that it can not wait for (OS) event(s).
	Extended Task A Task which can block by itself and wait for (OS) event(s).
Time Frame	The minimum inter-arrival time for a Task/ISR.
Trusted Function	A service provided by a trusted OS-Application that can be used by other OS-Applications (trusted or non-trusted).
Worst case execution time (WCET)	The longest possible execution time.
Write access	Storing a value in a register or memory location. All memory accesses that have the consequence of writing (e.g. reads that have the side effect of writing to a memory location) are treated as write accesses.

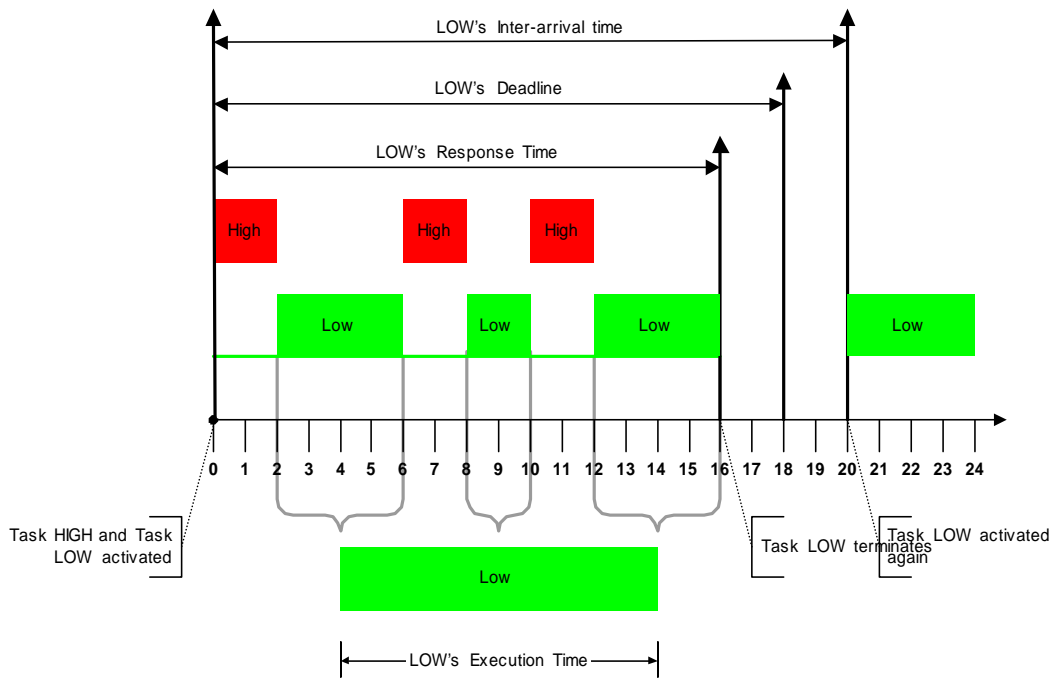


Figure 2.1: Definition of Timing Terminology

3 Related documentation

3.1 Input documents

- [1] Layered Software Architecture
AUTOSAR_EXP_LayeredSoftwareArchitecture.pdf
- [2] Requirements on Operating System
AUTOSAR_SRS_OS.pdf
- [3] General Requirements on Basic Software Modules
AUTOSAR_SRS_BSWGeneral.pdf
- [4] Specification of the Virtual Functional Bus
AUTOSAR_EXP_VFB.pdf
- [5] Requirements on Software FreeRunningTimer
AUTOSAR_SRS_FreeRunningTimer.pdf
- [6] Specification of GPT Driver
AUTOSAR_SWS_GPTDriver.pdf
- [7] Specification of Standard Types
AUTOSAR_SWS_StandardTypes.pdf
- [8] Specification of Memory Mapping
AUTOSAR_SWS_MemoryMapping.pdf
- [9] Specification of RTE
AUTOSAR_SWS_RTE.pdf
- [10] Specification of ECU Configuration
AUTOSAR_TPS_ECUConfiguration.pdf
- [11] Basic Software Module Description Template,
AUTOSAR_TPS_BSWModuleDescriptionTemplate.pdf
- [12] Specification of Multi-Core OS Architecture,
AUTOSAR_SWS_MultiCoreOS.pdf
- [13] List of Basic Software Modules,
AUTOSAR_TR_BSWModuleList.pdf
- [14] Specification of RTE,
AUTOSAR_SWS_RTE.pdf

3.2 Related standards and norms

3.2.1 OSEK/VDX

The OSEK/VDX specifications are publicly available from www.osek-vdx.org

- [15] Operating System
Version 2.2.3
17th February 2005
- [16] Time-Triggered Operating System
Version 1.0
24th July 2001
- [17] System Generation OIL: OSEK Implementation Language
Version 2.5
1st July 2004
- [18] OSEK RunTime Interface (ORTI) Part A: Language Specification
Version 2.2
14th November 2005
- [19] OSEK Run Time Interface (ORTI) Part B: OSEK Objects and Attributes
Version 2.2
25th November 2005
- [20] Binding Specification
Version 1.4.2
15th July 2004

3.2.2 HIS

The HIS (Hersteller Initiative Software) documents are publicly available from www.automotive-his.de

- [21] Requirements for Protected Applications under OSEK
Version 1
25th September 2002.
- [22] OSEK OS Extensions for Protected Applications
Version 1.0
27th July 2003

3.2.3 ISO/IEC

[23] ISO/IEC 9899:1990 Programming Language – C

(Remark: The international ISO standard ISO/IEC 9899:1990, also sometimes simply called »C90«, describes the language C. It was introduced in 1990 and replaced the ANSI C standard that was introduced only one year before, that's why it is also called »C89«. C89 differs from ISO/IEC 9899:1990 essentially only by the copyright note.)

[24] ISO/IEC 9899:1999 Programming Language – C

(Remark: A revised version of the standard was published in 1999. It is officially ISO/IEC 9899:1999, but is more often referred to as »C99«.)

3.3 Company Reports, Academic Work, etc.

[25] Extensions of OSEK OS for Protected Applications
OSEK Support Project DC058_02
DaimlerChrysler AG

4 Constraints and assumptions

4.1 Existing Standards

This document makes the following assumptions about the referenced related standards and norms:

- OSEK OS [15] provides a sufficiently flexible scheduling policy to schedule AUTOSAR systems.
- OSEK OS [15] is a mature specification and implementations are used in millions of ECUs worldwide.
- OSEK OS [15] does not provide sufficient support for isolating multi-source software components at runtime.
- OSEK OS [15] does not provide sufficient runtime support for demonstrating the absence of some classes of fault propagation in a safety-case.
- OSEKtime OS [16] and the HIS Protected OSEK [22] are immature specifications that contain concepts necessary for AUTOSAR and satisfy specific application domains. It is the purpose of this document to identify these needs and to recommend the use of parts (or all) of these specifications as appropriate.

4.2 Terminology

The specification uses the following operators when requirements specify multiple terms:

NOT : negation of a single term e.g. NOT Weekend

AND : conjunction of two terms e.g. Weekend AND Saturday

OR : disjunction of two terms e.g. Monday OR Tuesday

A requirement comprising multiple terms is evaluated left to right.

The precedence rules are:

Highest Precedence	NOT
Lowest Precedence	AND OR

The expression NOT X AND Y means (NOT X) AND (Y)

Where operators of the same precedence are used in the same sentence, commas are used to disambiguate. The expression X AND Y, OR Z means (X AND Y) OR Z.

4.3 Interaction with the RTE

The configuration of an AUTOSAR system [4] maps the »runnables« of a »software component« to (one or more) tasks that are scheduled by the operating system. All runnables in a task share the same protection boundary. In AUTOSAR, a software component must not include an interrupt handler. A software component is therefore implemented as runnables executing within the body of a task, or set of tasks, only.

Runnables get access to hardware-sourced data through the AUTOSAR RTE. The RTE provides the runtime interface between runnables and the basic software modules. The basic software modules also comprise a number of tasks and ISRs that are scheduled by the operating system.

It is assumed that the software component templates and the description of the basic software modules provide sufficient information about the required runtime behavior to be able to specify the attributes of tasks required to configure the OS.

4.4 Operating System Abstraction Layer (OSAL)

Systems that do not use the OS defined in AUTOSAR can provide a platform for the execution of AUTOSAR software components using an Operating System Abstraction Layer. The interface to the OSAL is exactly that defined for the AUTOSAR OS.

4.5 Multi-Core Hardware assumptions

There are currently several existing and suggested HW-architectures¹ for Multi-Core microprocessors. There is considerable variation in the features offered by these architectures. Therefore this section attempts to capture a common set of architectural features required for Multi-Core.

Hardware assumptions shall remain assumptions and shall not become official Autosar requirements.

4.5.1 CPU Core features

1. More than one core on the same piece of silicon.
2. The HW offers a method that can be used by the SW to identify a core.
3. The hardware supports atomic read and atomic write operations for a fixed word length depending on the hardware.
4. The hardware supports some atomic Test-And-Set functionality or similar functionalities that can be used to built a critical section shared between cores. Additional atomic operations may exist.
5. The cores may have the same instruction set; at least a common basic instruction set is available on all cores. Core specific add-ons may exist but they are not taken into account.
6. The cores have the same data representation. For example, the same size of integer, same byte and bit order, etc.

¹ In this context "architecture" encompasses: the connections between cores and memory, and to peripherals and how interrupts work.

7. If per-core caches exist, AUTOSAR requires support for RAM - cache coherency in HW or in SW. In software means that the cache-controller can be programmed by the SW in a way that it invalidates cache lines or excludes certain memory regions from caching.
8. In case of an exception (such as an illegal memory reference or divide by zero) the exception occurs on the core that introduced the exception.
9. For notification purposes, it is possible to trigger an interrupt/trap on any core.

4.5.2 Memory features

1. Shared RAM is available to all cores; at least all cores can share a substantial part of the memory.
2. Flash shall be shared between all cores at least. However, performance can be improved if Flash/RAM can be partitioned so that there are separate pathways from cores to Flash.
3. A single address space is assumed, at least in the shared parts of the memory address space.
4. The AUTOSAR Multi-Core architecture shall be capable to run on systems that do and do not support memory protection. If memory protection exists, all cores are covered by a hardware based memory protection.

4.5.3 Multi-Core Limitations

- In AUTOSAR R4.0, it is not supported to activate additional cores under control of AUTOSAR after the Operating System was started.
- The scheduling algorithm does not assign TASKs dynamically to cores.
- The AUTOSAR OS RESOURCE algorithm is not supported across cores. RESOURCES can be used locally, between TASKs that are bound to the same core but not between TASKs/ISRs which are bound to different cores.

4.6 Limitations

4.6.1 Hardware

The core AUTOSAR operating system assumes free access to hardware resources, which are managed by the OS itself. This includes, but is not limited to, the following hardware:

- interrupt control registers
- processor status words
- stack pointer(s)

Specific (extended) features of the core operating system extend the requirements on hardware resource. The following list outlines the features that have requirements on the hardware. Systems that do not use these OS features do not have these hardware requirements.

- **Memory Protection:** A hardware memory protection unit is required. All memory accesses that have the consequence of writing (e.g. reads that have the side effect of writing to a memory location) shall be treated as writes.
- **Time Protection:** Timer Hardware for monitoring execution times and arrival rates.
- **»Privileged« and »non-privileged« modes on the MCU:** to protect the OS against internal corruption caused by writes to OS controlled registers. This mode must not allow OS-Applications to circumvent protection (e.g. write registers which govern memory protection, write to processor status word etc.). The privileged mode must be under full control of the protected OS which uses the mode internally and to transfer control back and forth from a non-trusted OS-Application to a trusted OS-Application. The microprocessor must support a controlled means which moves a processor into this privileged mode.
- **Local/Global Time Synchronization:** A global time source is needed.

In general hardware failures in the processor are not detected by the operating system. In the event of hardware failure, correct operation of the OS cannot be guaranteed.

The resources managed by a specific OS implementation have to be defined within the appropriate configuration file of the OS.

4.6.2 Programming Language

The API of the operating system is defined as C89 [23] function calls or macros. If other languages are used they must adapt to the C interface. This is because C99 [24] allows for internal dynamic memory allocation during subroutine calls. Most automotive applications are static (non-heap based).

4.6.3 Miscellaneous

The operating system does not provide services for dynamic memory management.

This specification handles only single core MCUs (this means one single thread of execution at one time). If you need support for multi-core MCUs please see [12].

If you are using a multi processor system and want to use this operating system, each processor has to run its own operating system. (e.g. a multi-processor system must use a different OS image for each processor.

4.7 Applicability to car domains

The operating system has the same design constraints regarding size and scalability under which the OSEK OS was designed. The immediate domain of applicability is therefore currently body, chassis and power train ECUs. However, there is no reason that the OS cannot be used to implement ECUs for infotainment applications.

5 Dependencies to other modules

There are no forced dependencies on other modules, however:

- It is assumed that the operating system may use timer units directly to drive counters.
- If the user needs to drive scheduling directly from global time, then a global time interrupt is required.
- If the user needs to synchronize the processing of a schedule table to a global time, the operating system needs to be told the global time using the `SyncScheduleTable()` service.
- The IOC described in this document provides communication between OS-Applications. The IOC generation is based on configuration information which is generated by the RTE generator. On the other hand the RTE uses functions generated by the IOC to transmit data.

5.1 File structure

5.1.1 Code file structure

The code file structure of the Operating system module is not fixed, besides the requirements in the General SRS.

5.1.2 Header file structure

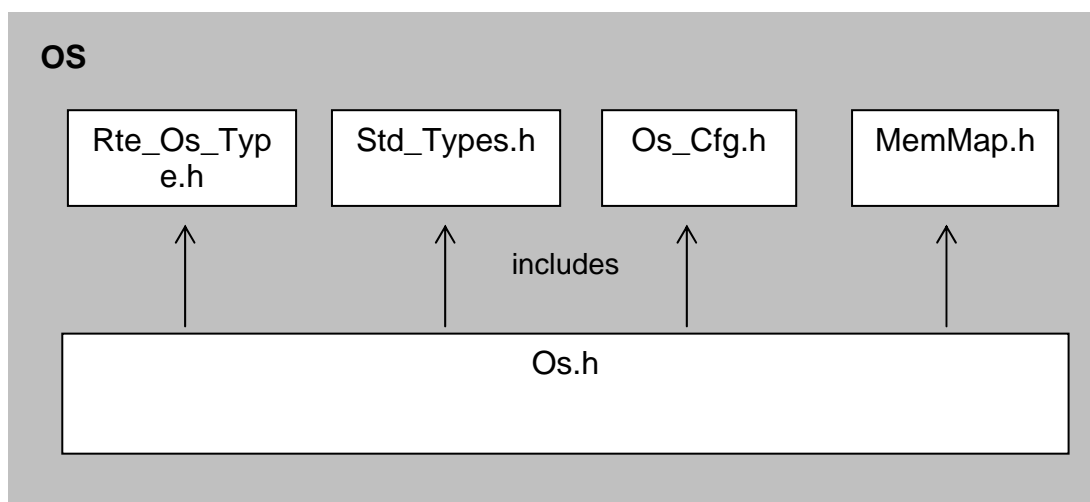


Figure 5:1: Header File Structure for the OS

The figure above contains the defined AUTOSAR header file hierarchy of the Operating System module.

The IOC generator generates an additional header file `loc.h`. Users of the `loc.h` shall include the `loc.h` file. If an implementation of the IOC requires additional header files, it is free to include them. The header files are self-contained, that means they will include all other header files, which they require.

[OS546] [OS implementer shall provide a header file structure, so that users of the Operating System module needs only to include the **Os.h** file] ()

[OS765] [The module header file `Os.h` shall include `Rte_Os_Type.h` to include the types which are common used by BSW Modules and Software Components. This file shall only contain types, that are not already defined in `Rte_Os_Type.h`.] ()

If an implementation of the Operating System module requires additional header files, it is free to include them. The header files are self contained, that means they will include all other header files which are required by them.

[OS552] [The Operating System module shall avoid the integration of incompatible (c or h) files by the following pre-processor checks:

For included (external) header files:

- `<MODULENAME>_AR_RELEASE_MAJOR_VERSION`
- `<MODULENAME>_AR_RELEASE_MINOR_VERSION`

shall be verified.

If the values are not identical to the values expected by the Operating System module, an error shall be reported.] ()

6 Requirements Traceability

This chapter contains references to requirements of other AUTOSAR documents.

Requirement	Satisfied by
-	OS388
-	OS439
-	OS274
-	OS265
-	OS524
-	OS457
-	OS391
-	OS764
-	OS422
-	OS510
-	OS278
-	OS276
-	OS560
-	OS516
-	OS264
-	OS555
-	OS270
-	OS564
-	OS561
-	OS262
-	OS475
-	OS313
-	OS544
-	OS280
-	OS330
-	OS462
-	OS533
-	OS425
-	OS281
-	OS344
-	OS493
-	OS511
-	OS258
-	OS408
-	OS531

-	OS443
-	OS436
-	OS321
-	OS402
-	OS499
-	OS466
-	OS455
-	OS267
-	OS416
-	OS211
-	OS179
-	OS266
-	OS410
-	OS506
-	OS282
-	OS399
-	OS172
-	OS348
-	OS460
-	OS279
-	OS459
-	OS562
-	OS085
-	OS349
-	OS529
-	OS209
-	OS502
-	OS536
-	OS543
-	OS355
-	OS275
-	OS320
-	OS271
-	OS445
-	OS505
-	OS323
-	OS565
-	OS273
-	OS556
-	OS403

-	OS541
-	OS450
-	OS515
-	OS409
-	OS440
-	OS525
-	OS563
-	OS347
-	OS324
-	OS437
-	OS549
-	OS762
-	OS501
-	OS454
-	OS763
-	OS356
-	OS539
-	OS535
-	OS507
-	OS449
-	OS467
-	OS520
-	OS269
-	OS503
-	OS552
-	OS446
-	OS537
-	OS418
-	OS514
-	OS412
-	OS547
-	OS362
-	OS111
-	OS006
-	OS532
-	OS381
-	OS463
-	OS500
-	OS495
-	OS397

-	OS283
-	OS396
-	OS447
-	OS442
-	OS017
-	OS444
-	OS415
-	OS268
-	OS292
-	OS542
-	OS236
-	OS054
-	OS353
-	OS376
-	OS548
-	OS272
-	OS314
-	OS522
-	OS453
-	OS518
-	OS177
-	OS016
-	OS058
-	OS404
-	OS438
-	OS293
-	OS343
-	OS083
-	OS060
-	OS765
-	OS421
-	OS194
-	OS389
-	OS367
-	OS198
-	OS242
-	OS513
-	OS423
-	OS521
-	OS430

-	OS519
-	OS504
-	OS526
-	OS287
-	OS494
-	OS027
-	OS358
-	OS546
-	OS277
-	OS496
-	OS458
-	OS354
-	OS290
-	OS308
-	OS263
-	OS483
-	OS050
-	OS509
-	OS112
-	OS328
-	OS369
-	OS557
-	OS414
-	OS291
-	OS411
-	OS538
-	OS427
-	OS350
-	OS419
-	OS368
-	OS550
-	OS071
-	OS429
-	OS424
-	OS452
-	OS530
-	OS540
-	OS484
-	OS226
-	OS387

-	OS304
-	OS559
-	OS332
-	OS261
-	OS237
-	OS097
-	OS498
-	OS551
-	OS401
-	OS407
-	OS009
-	OS456
-	OS420
-	OS566
-	OS417
-	OS413
-	OS527
-	OS431
-	OS464
-	OS364
-	OS512
-	OS284
-	OS508
-	OS256
-	OS365
-	OS435
-	OS461
-	OS289
-	OS351
-	OS285
-	OS361
-	OS534
-	OS451
-	OS497
-	OS554
-	OS173
-	OS553
-	OS545
-	OS303
-	OS312

-	OS300
-	OS528
-	OS517
-	OS523
-	OS476
-	OS225
-	OS239
-	OS448
-	OS428
-	OS309
-	OS100
-	OS045
-	OS327
-	OS311
BSW003	OS767
BSW00301	OS767
BSW00302	OS767
BSW00304	OS767
BSW00305	OS767
BSW00306	OS767
BSW00307	OS767
BSW00308	OS767
BSW00309	OS767
BSW00310	OS767
BSW00312	OS767
BSW00314	OS767
BSW00318	OS767
BSW00321	OS767
BSW00325	OS767
BSW00326	OS767
BSW00327	OS767
BSW00328	OS767
BSW00329	OS767
BSW00330	OS767
BSW00333	OS767
BSW00334	OS767
BSW00335	OS767
BSW00337	OS767
BSW00338	OS767
BSW00339	OS767

BSW00341	OS767
BSW00342	OS767
BSW00344	OS767
BSW00347	OS767
BSW00350	OS767
BSW00355	OS767
BSW00357	OS767
BSW00358	OS767
BSW00361	OS767
BSW00369	OS767
BSW00370	OS767
BSW00373	OS767
BSW00374	OS767
BSW00375	OS767
BSW00376	OS767
BSW00377	OS767
BSW00378	OS767
BSW00379	OS767
BSW00380	OS767
BSW00381	OS767
BSW00383	OS767
BSW00384	OS767
BSW00385	OS767
BSW00386	OS767
BSW00401	OS767
BSW00404	OS767
BSW00405	OS767
BSW00406	OS767
BSW00407	OS767
BSW00409	OS767
BSW00410	OS767
BSW00411	OS767
BSW00412	OS767
BSW00413	OS767
BSW00414	OS767
BSW00415	OS767
BSW00417	OS767
BSW00419	OS767
BSW00422	OS767
BSW00423	OS767

BSW00435	OS767
BSW00436	OS767
BSW00437	OS767
BSW00439	OS767
BSW00440	OS767
BSW00441	OS767
BSW006	OS767
BSW007	OS767
BSW009	OS767
BSW010	OS767
BSW097	OS001
BSW098	OS007, OS002
BSW099	OS191
BSW11000	OS026
BSW11001	OS056
BSW11002	OS199, OS227, OS013, OS201, OS206
BSW11003	OS067, OS068
BSW11005	OS195, OS208, OS207
BSW11006	OS196, OS087, OS086
BSW11007	OS081
BSW11008	OS469, OS048, OS465, OS473, OS474, OS471, OS472, OS470, OS028, OS033, OS037, OS089, OS064
BSW11009	OS051, OS052, OS088, OS093, OS092, OS069, OS070
BSW11010	OS056
BSW11011	OS096, OS245
BSW11012	OS241, OS240
BSW11013	OS044, OS051, OS056, OS033, OS037, OS088, OS093, OS068, OS064, OS070, OS246, OS210
BSW11014	OS033, OS037, OS110, OS244, OS243, OS106, OS109, OS107, OS108
BSW11016	OS241, OS240
BSW11018	OS299
BSW11019	OS336
BSW11020	OS286
BSW11021	OS301
BSW161	OS767
BSW162	OS767
BSW168	OS767
BSW170	OS767
BSW172	OS767
BSW4080001	OS600, OS606, OS616, OS628, OS627, OS626, OS568, OS569, OS674, OS673, OS672, OS675, OS579, OS583, OS596

BSW4080003	OS571, OS570, OS573
BSW4080005	OS667, OS571, OS570, OS573, OS572
BSW4080006	OS608, OS609, OS607, OS610, OS625, OS668, OS669, OS575, OS574, OS572, OS670, OS679, OS677, OS678, OS676, OS578, OS579, OS576, OS577, OS681, OS680, OS683, OS682, OS685, OS684, OS584, OS585, OS580, OS581, OS582
BSW4080007	OS616, OS617, OS621, OS586, OS588, OS716, OS587, OS713, OS715, OS714
BSW4080008	OS567, OS582
BSW4080011	OS583
BSW4080013	OS607, OS618, OS619, OS629, OS623, OS639, OS636, OS635, OS638, OS637, OS631, OS630, OS645, OS643, OS647, OS646, OS640, OS663, OS664, OS665, OS569, OS589, OS594, OS595, OS592, OS593, OS590, OS591
BSW4080015	OS600, OS596, OS599, OS598
BSW4080016	OS604, OS605, OS602
BSW4080018	OS632, OS634, OS633, OS644, OS642, OS649, OS648, OS641, OS654, OS653, OS656, OS655, OS658, OS657, OS659, OS650, OS652, OS661, OS660
BSW4080020	OS611, OS671, OS761, OS760, OS751, OS750, OS756, OS757, OS758, OS759, OS752, OS753, OS754, OS755, OS740, OS747, OS748, OS745, OS746, OS743, OS744, OS741, OS742, OS749, OS739, OS738, OS731, OS730, OS733, OS732, OS735, OS734, OS737, OS736, OS729, OS728, OS727, OS722, OS721, OS720, OS726, OS725, OS724, OS723, OS719, OS718
BSW4080021	OS612, OS613, OS614, OS615, OS624, OS622, OS620, OS649, OS648, OS654, OS653, OS656, OS655, OS658, OS657, OS659, OS650, OS651, OS652, OS661, OS660, OS662, OS666, OS686, OS687, OS688, OS689, OS692, OS691, OS690, OS696, OS695, OS694, OS693, OS699, OS697, OS698, OS712, OS711, OS710, OS709, OS708, OS707, OS706, OS705, OS704, OS703, OS702, OS701, OS700
BSW4080026	OS575, OS574, OS679, OS677, OS678, OS676, OS576, OS577, OS681, OS680, OS683, OS682, OS685, OS684, OS584, OS585
BSW4080027	OS575, OS574, OS679, OS677, OS678, OS676, OS576, OS577, OS681, OS680, OS683, OS682, OS685, OS684, OS584, OS585
SWFRT00020	OS374
SWFRT00022	OS370
SWFRT00025	OS392, OS383
SWFRT00030	OS384
SWFRT00031	OS384
SWFRT00032	OS767
SWFRT00033	OS377
SWFRT00034	OS382
SWFRT00047	OS393

7 Functional specification

7.1 Core OS

7.1.1 Background & Rationale

The OSEK/VDX Operating System [15] is widely used in the automotive industry and has been proven in use in all classes of ECUs found in modern vehicles. The concepts that OSEK OS has introduced are widely understood and the automotive industry has many years of collective experience in engineering OSEK OS based systems.

OSEK OS is an event-triggered operating system. This provides high flexibility in the design and maintenance of AUTOSAR based systems. Event triggering gives freedom for the selection of the events to drive scheduling at runtime, for example angular rotation, local time source, global time source, error occurrence etc.

For these reasons the core functionality of the AUTOSAR OS shall be based upon the OSEK OS. In particular OSEK OS provides the following features to support concepts in AUTOSAR:

- fixed priority-based scheduling
- facilities for handling interrupts
- only interrupts with higher priority than tasks
- some protection against incorrect use of OS services
- a startup interface through `StartOS()` and the `StartupHook()`
- a shutdown interface through `ShutdownOS()` and the `ShutdownHook()`

OSEK OS provides many features in addition to these. Readers should consult the OSEK specification [15] for details.

Basing AUTOSAR OS on OSEK OS means that legacy applications will be backward compatible – i.e. applications written for OSEK OS will run on AUTOSAR OS. However, some of the features introduced by AUTOSAR OS require restrictions on the use of existing OSEK OS features or extend existing OSEK OS features.

7.1.2 Requirements

[OS001] [The Operating System module shall provide an API that is backward compatible with the OSEK OS API [15].] (BSW097)

7.1.2.1 Restrictions on OSEK OS

It is too inefficient to achieve timing and memory protection for alarm callbacks. They are therefore not allowed in specific scalability classes ([OS242](#))

[OS242] [The Operating System module shall only allow Alarm Callbacks in Scalability Class 1.] ()

OSEK OS is required to provide functionality to handle inter-task (internal) communication according to the OSEK COM specification when internal communication only is required in the system. In AUTOSAR, internal communication is provided by the AUTOSAR RTE or by AUTOSAR COM at least one of which will be present for all AUTOSAR ECUs.

AUTOSAR OS, when used in an AUTOSAR system, therefore does not need to support internal communication.

An OSEK OS must implement internal communication if the symbol `LOCALMESSAGEONLY` is defined. AUTOSAR OS can deprecate the need to implement OSEK COM functionality and maintain compatibility with OSEK suite of specifications by ensuring that AUTOSAR OS always exists in an environment where `LOCALMESSAGEONLY` is undefined.

OSEK OS has one special resource called `RES_SCHEDULER`. This resource has 2 specific aspects:

1. It is always present in the system, even if it is not configured. This means that the `RES_SCHEDULER` is always known by the OS.
2. It has always the highest Task priority. This means a Task which allocates this resource can not be preempted by other Tasks.

Since special cases are always hard to handle (e.g. in this case with respect to timing protection) AUTOSAR OS handles `RES_SCHEDULER` as any other resource. This means that the `RES_SCHEDULER` is not automatically created. However, a configuration attribute allows that a resource in AUTOSAR OS can optionally be assigned the priority of the highest priority task in the system.

For backwards compatibility with OSEK OS systems, see Chapter 12.8 on how to configure a standard resource called `RES_SCHEDULER` in a way that make it compatible with the resource of the same name which is declared automatically in OSEK OS.

In OSEK OS users must declare Operating System objects with specific macros (e.g. `DeclareTask()`, ...) An AUTOSAR OS implementation shall not depend on such declarations and shall (for backwards compatibility) supply macros without functionality.

7.1.2.2 Undefined Behaviour in OSEK OS

There are a number of cases where the behaviour of OSEK OS is undefined. These cases represent a barrier to portability. AUTOSAR OS tightens the OSEK OS specification by defining the required behaviour.

[OS304] [If in a call to `SetRelAlarm()` the parameter “increment” is set to zero, the service shall return `E_OS_VALUE` in standard and extended status .] ()

[OS424] [The first call to `startOS()` (for starting the Operating System) shall not return.] ()

[OS425] [If `shutdownOS()` is called and `ShutdownHook()` returns then the Operating System module shall disable all interrupts and enter an endless loop.] ()

7.1.2.3 Extensions to OSEK OS

[OS299] [The Operating System module shall provide the services `DisableAllInterrupts()`, `EnableAllInterrupts()`, `SuspendAllInterrupts()`, `ResumeAllInterrupts()` prior to calling `startOS()` and after calling `shutdownOS()`.] (BSW11018)

It is assumed that the static variables of the functions mentioned in [OS299](#) are initialized.

[OS301] [The Operating System module shall provide the ability to increment a software counter as an alternative action on alarm expiry.] (BSW11021)

The Operating System module provides API service `IncrementCounter()` (see [OS399](#)) to increment a software counter.

[OS476] [The Operating System module shall allow to automatically start preconfigured absolute alarms during the start of the Operating System.] ()

[OS476](#) is an extension to OSEK OS which allows this only for relative alarms.

[OS566] [The Operating System API shall check in extended mode all pointer arguments for a `NULL` pointer and return `OS_E_PARAM_POINTER` if such an argument is `NULL`.] ()

7.2 Software Free Running Timer

Due to the fact that the number of timers is often very limited, some functionality and configuration is added to extend the reuse of timers. E.g. this allows timer measurements. For more details see also [5] (SWFRT).

[OS374] [The Operating System module shall handle all the initialization and configuration of timers used directly by the Operating System module and not handled by the GPT driver.] (SWFRT00020)

The Operating System module provides API service `GetCounterValue()` (see [OS383](#)) to read the current count value of a counter (returning either the hardware timer ticks if counter is driven by hardware or the software ticks when user drives counter).

The Operating System module provides API service `GetElapsedValue()` (see [OS392](#)) to get the number of ticks between the current tick value and a previously read tick value.

[OS384] [The Operating System module shall adjust the read out values of hardware timers (which drive counters) in such that the lowest value is zero and consecutive reads return an increasing count value until the timer wraps at its modulus.] (SWFRT00030, SWFRT00031)

7.3 Schedule Tables

7.3.1 Background & Rationale

It is possible to implement a statically defined task activation mechanism using an OSEK counter and a series of auto started alarms. In the simple case, this can be achieved by specifying that the alarms are not modified once started. Run-time modifications can only be made if relative synchronization between alarms can be guaranteed. This typically means modifying the alarms while associated counter tick interrupts are disabled.

Schedule Tables address the synchronization issue by providing an encapsulation of a statically defined set of expiry points. Each expiry point defines:

- one or more actions that must occur when it is processed where an action is the activation of a task or the setting of an event.
- An offset in ticks from the start of the schedule table

Each schedule table has a duration in ticks. The duration is measured from zero and defines the modulus of the schedule table.

At runtime, the Operating System module will iterate over the schedule table, processing each expiry point in turn. The iteration is driven by an OSEK counter. It therefore follows that the properties of the counter have an impact on what is possible to configure on the schedule table.

7.3.2 Requirements

7.3.2.1 Structure of a Schedule Table

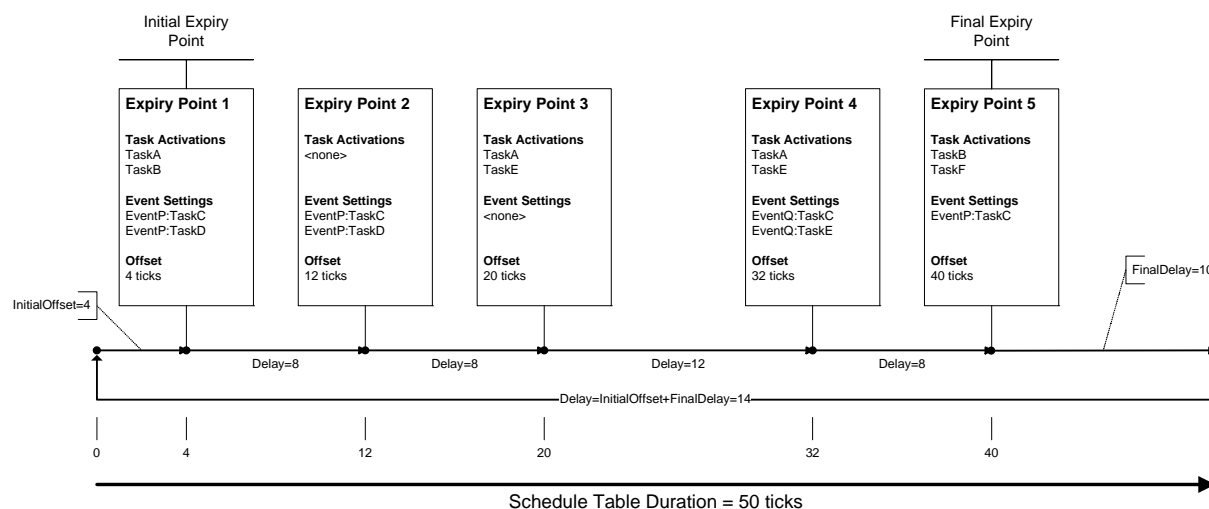


Figure 7.1: Anatomy of a Schedule Table

[OS401] [A schedule table shall have at least one expiry point.] ()

[OS402] [An expiry point shall contain a (possibly empty) set of tasks to activate.] ()

[OS403] [An expiry point shall contain a (possibly empty) set of events to set.] ()

[OS404] [An expiry point shall contain an offset in ticks from the start of the schedule table.] ()

7.3.2.2 Constraints on Expiry Points

There is no use case for an empty expiry point, so each one must define at least one action.

[OS407] [An expiry point shall activate at least one task OR set at least one event.] ()

The OS needs to know the order in which expiry points are processed. It is therefore necessary to ensure that the expiry points on a schedule table can be totally ordered. This is guaranteed by forcing each expiry point on a schedule table to have a unique offset.

[OS442] : [Each expiry point on a given schedule table shall have a unique offset.] ()

Iteration over expiry points on a schedule table is driven by an OSEK counter. The characteristics of the counter – `OsCounterMinCycle` and `OsCounterMaxAllowedValue` – place constraints on expiry point offsets.

[OS443] [The Initial Offset shall be zero OR in the range `OsCounterMinCycle` .. `OsCounterMaxAllowedValue` of the underlying counter.] ()

Similarly, constraints apply to the delays between of adjacent expiry points and the delay to the logical end of the schedule table.

[OS408] [The delay between adjacent expiry points shall be in the range `OsCounterMinCycle` .. `OsCounterMaxAllowedValue` of the underlying counter.] ()

7.3.2.3 Processing Schedule Tables

[OS002] [The Operating System module shall process each expiry point on a schedule table from the Initial Expiry Point to the Final Expiry Point in order of increasing offset.] (BSW098)

[OS007] [The Operating System module shall permit multiple schedule tables to be processed concurrently.] (BSW098)

[OS409] [A schedule table of the Operating System module shall be driven by exactly one counter.] ()

[OS410] [The Operating System module shall be able to process at least one schedule table per counter at any given time.] ()

[OS411] [The Operating System module shall make use of ticks so that one tick on the counter corresponds to one tick on the schedule table.] ()

It is possible to activate a task and set (one or more unique) events for the same task at the same expiry point. The ordering of task activations and event settings performed from the expiry point could lead to different implementations exhibiting different behaviour (for example, activating a suspended task and then setting an event on the task would succeed but if the ordering was reversed then the event setting would fail). To prevent such non-determinism, it is necessary to enforce a strict ordering of actions on the expiry point.

[OS412] [The Operating System module shall process all task activations on an expiry point first and then set events.] ()

A schedule table always has a defined state and the following figure illustrates the different states (for a non-synchronized schedule table) and the transitions between them.

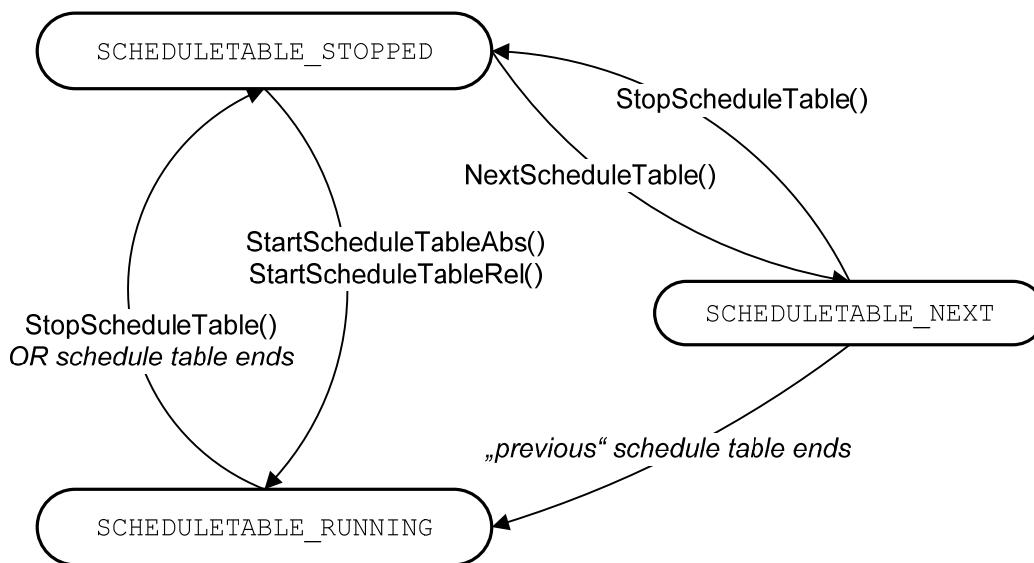


Figure 7.2: States of a schedule table

If a schedule table is not active – this means that is not processed by the Operating System – the state is SCHEDULETABLE_STOPPED. After starting a schedule tables enters the SCHEDULETABLE_RUNNING state where the OS processes the expiry points. If the service to switch a schedule table is called a schedule table enters the the SCHEDULETABLE_NEXT state and waits until the “current” schedule table ends.

7.3.2.4 Repeated Schedule Table Processing

A schedule table may or may not repeat after the final expiry point is processed. This allows two types of behaviour:

1. single-shot – the schedule table processes each expiry point in sequence and then stops at the end. This is useful for triggering a phased sequence of actions in response to some trigger
2. repeating – the schedule table processes each expiry point in turn, After processing the final expiry point, it loops back to the initial expirt point. This is useful for building applications that perform repeated processing or system which need to synchronise processing to a driver source.

A repeating schedule table means that each expiry point is repeated at a period equal to the schedule table duration.

[OS413] [The schedule table shall be configurable as either single-shot or repeating.] ()

[OS009] [If the schedule table is single-shot, the Operating System module shall stop the processing of the schedule table Final Delay ticks after the Final Expiry Point is processed.] ()

[OS427] [If the schedule table is single-shot, the Operating System module shall allow a Final Delay between 0 .. `OsCounterMaxAllowedValue` of the underlying counter.] ()

[OS444] [For periodic schedule tables the value of Final Delay shall be in the range `OsCounterMinCycle` .. `OsCounterMaxAllowedValue` of the underlying counter.] ()

[OS194] [After processing the Final Expiry Point, and if the schedule table is repeating, the Operating System shall process the next Initial Expiry Point, after Final Delay plus Initial Offset ticks have elapsed.] ()

7.3.2.5 Controlling Schedule Table Processing

The application is responsible for starting and stopping the processing of a schedule table.

The Operating System module provides the service `StartScheduleTableAbs()` (see [OS358](#)) to start the processing of a schedule table at an absolute value “Start” on the underlying counter. (The Initial Expiry Point has to be processed when the value of the underlying counter equals `Start + InitialOffset`).

The Operating System module provides the service `StartScheduleTableRel()` (see [OS347](#)) to start the processing of a schedule table at “Offset” relative to the “Now” value on the underlying counter (The Initial Expiry Point shall be processed when the value of the underlying counter equals `Now + Offset + InitialOffset`).

The figure below illustrates the two different methods for a schedule table driven by a counter with a modulus of 65536 (i.e. an `OsCounterMaxAllowedValue = 65535`).

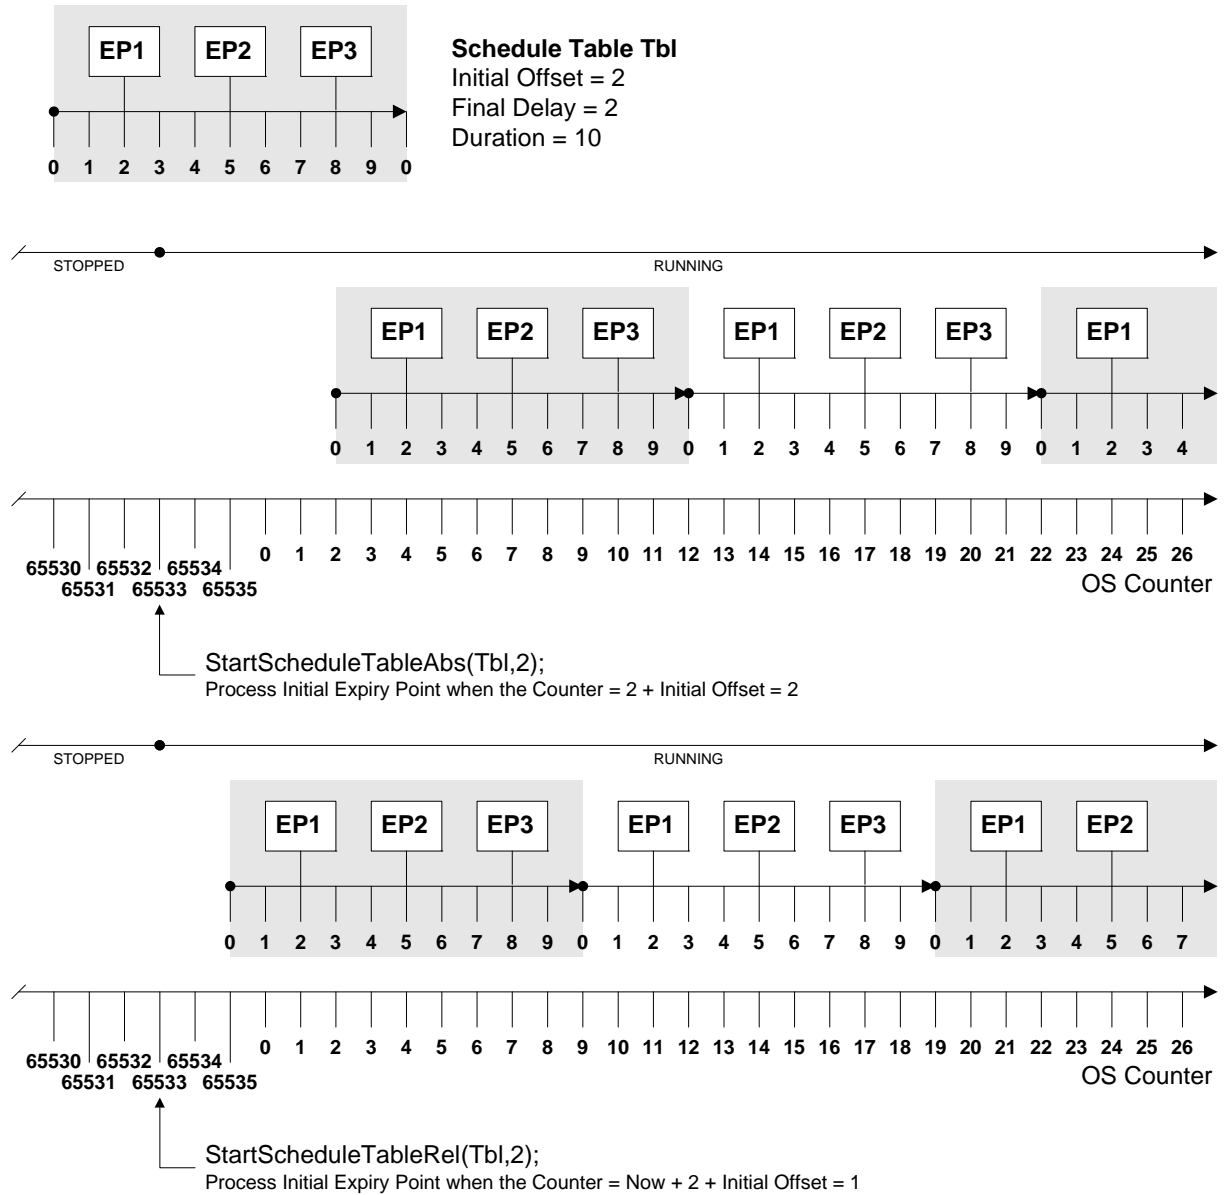


Figure 7.3: Starting a Schedule Table at an Absolute and a Relative Count

The Operating System module provides the service `StopScheduleTable()` (see [OS006](#)) to cancel the processing of a schedule table immediately at any point while the schedule table is running.

[OS428] [If schedule table processing has been cancelled before reaching the Final Expiry Point and is subsequently restarted then [OS358/OS347](#) means that the re-start occurs from the start of the schedule table.] ()

The Operating System module provides the service `NextScheduleTable()` (see [OS191](#)) to switch the processing from one schedule table to another schedule table.

[OS414] [When a schedule table switch is requested, the OS shall continue to process expiry points on the current schedule table. After the Final Expiry Point there will be a delay equivalent to Final Delay ticks before processing the switched-to schedule table. The initial expiry point will be processed after initial offset.] ()

The Operating System module provides the service `GetScheduleTableStatus()` (see [OS227](#)) to query the state of a schedule table.

Schedule tables can be configured (see chapter 10) to start automatically during start of the Operating System module (like Tasks and Alarms in OSEK OS). OSEK OS defines a specific order: Autostart of Tasks is performed before autostart of alarms. AUTOSAR OS extends this with schedule tables.

[OS510] [The Operating System module shall perform the autostart of schedule tables during startup after the autostart of Tasks and Alarms.] ()

7.4 Schedule Table Synchronization

7.4.1 Background & Rationale

The absolute time at which the Initial Expiry Point on a schedule table is processed is under user control. However, if the schedule table repeats then it is not guaranteed that the absolute count value at which the initial expiry point was first processed is the same count value at which it is subsequently processed. This is because the duration of the schedule table need not be equal to the counter modulus.

In many cases it may be important that schedule table expiry points are processed at specific absolute values of the underlying counter. This is called **synchronization**. Typical use-cases include:

- Synchronization of expiry points to degrees of angular rotation for motor management
- Synchronizing the computation to a global (network) time base. Note that in AUTOSAR, the Operating System does not provide a global (network) time source because
 1. a global time may not be needed in many cases
 2. other AUTOSAR modules, most notably FlexRay, provide this independently to the Operating System
 3. if the Operating System is required to synchronize to multiple global (network) time sources (for example when building a gateway between two time-triggered networks) the Operating System cannot be the source of a unique global time.

AUTOSAR OS provides support for synchronization in two ways:

1. implicit synchronization – the counter driving the schedule table is the counter with which synchronization is required. This is typically how synchronization with time-triggered networking technologies (e.g. FlexRay, TTP) is achieved – the underlying hardware manages network time synchronization and simply presents time as an output/compare timer interface to the Operating System. The following figure shows the possible states for schedule tables with implicit synchronization.

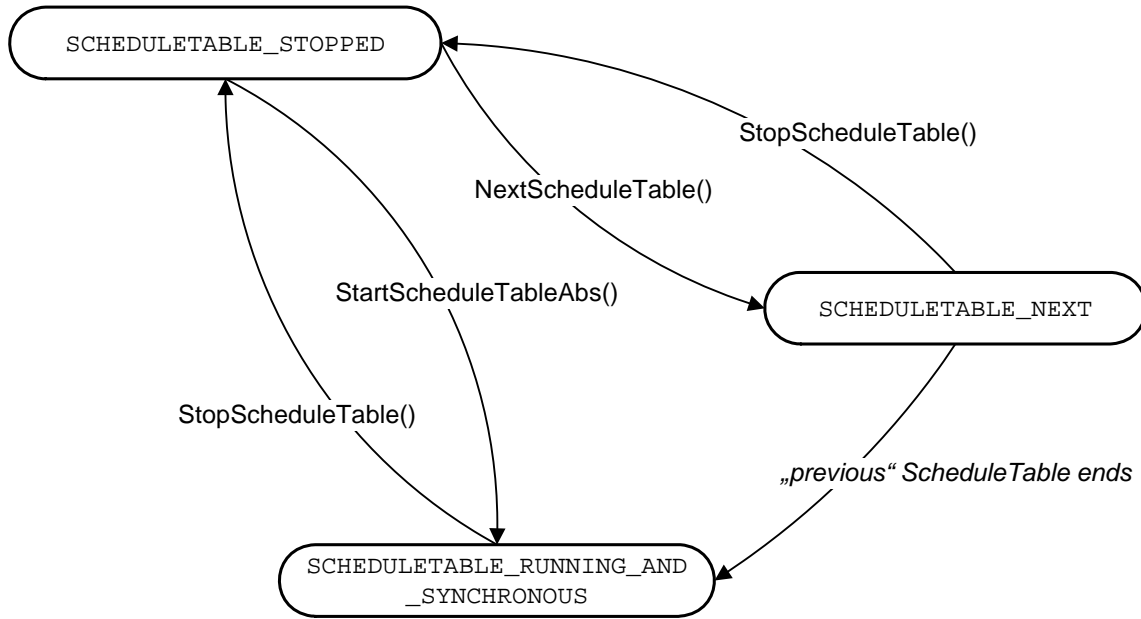


Figure 7.4: States of an implicit synchronized schedule table

2. explicit synchronization – the schedule table is driven by an Operating System counter which is **not** the counter with which synchronization is required. The Operating System provides additional functionality to keep schedule table processing driven by the Operating System counter synchronized with the synchronization counter. This is typically how synchronization with periodically broadcast global times works. The next figure shows the states of such schedule tables.

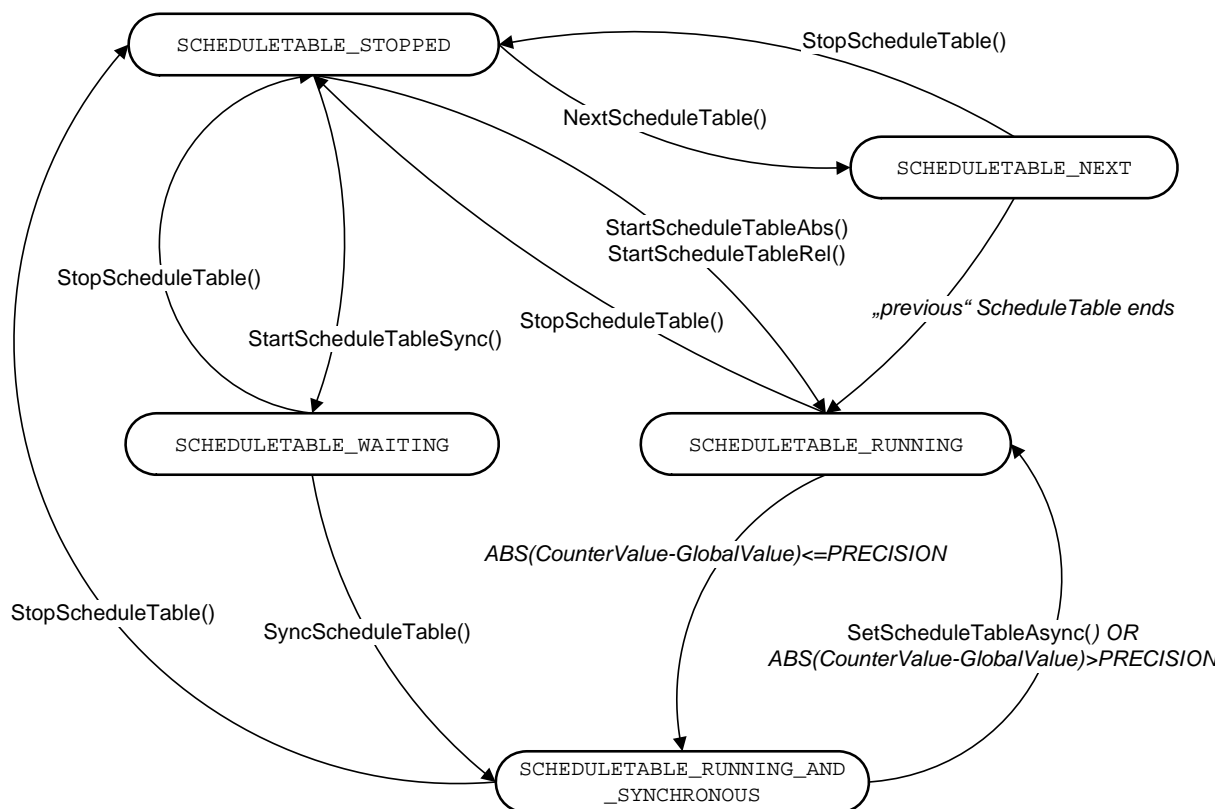


Figure 7.5: States of an explicit synchronized schedule table (not all conditions for transitions are shown in the picture)

7.4.2 Requirements

[OS013] [The Operating System module shall provide the ability to synchronize the processing of schedule table to known counter values.] (BSW11002)

7.4.2.1 Implicit Synchronization

The Operating System module does not need to provide any additional support for implicit synchronization of schedule tables. However, it is necessary to constrain configuration and runtime control of the schedule table so that ticks on the configured schedule table can be aligned with ticks on the counter. This requires the range of the schedule table to be identical to the range of the counter (the equality of tick resolution of each is guaranteed by the requirements on the schedule table / counter interaction):

[OS429] [A schedule table of the Operating System module that is implicitly synchronized shall have a Duration equal to `OsCounterMaxAllowedValue + 1` of its associated OSEK OS counter.] ()

To synchronize the processing of the schedule table it must be started at a known counter value. The implication of this is that a schedule table requiring implicit

synchronization must only be started at an absolute counter value and cannot be started at a relative count value.

[OS430] [The Operating System module shall prevent a schedule table that is implicitly synchronized from being started at a relative count value.] ()

When the schedule table is started at an absolute counter value each expiry point will be processed when the counter equals the value specified in the service call plus expiry point's offset. The common use-case is to ensure that the offsets specified in the schedule table configuration correspond to absolute values of the underlying counter. This is achieved trivially using `StartScheduleTable(Tbl,0)` as shown below.

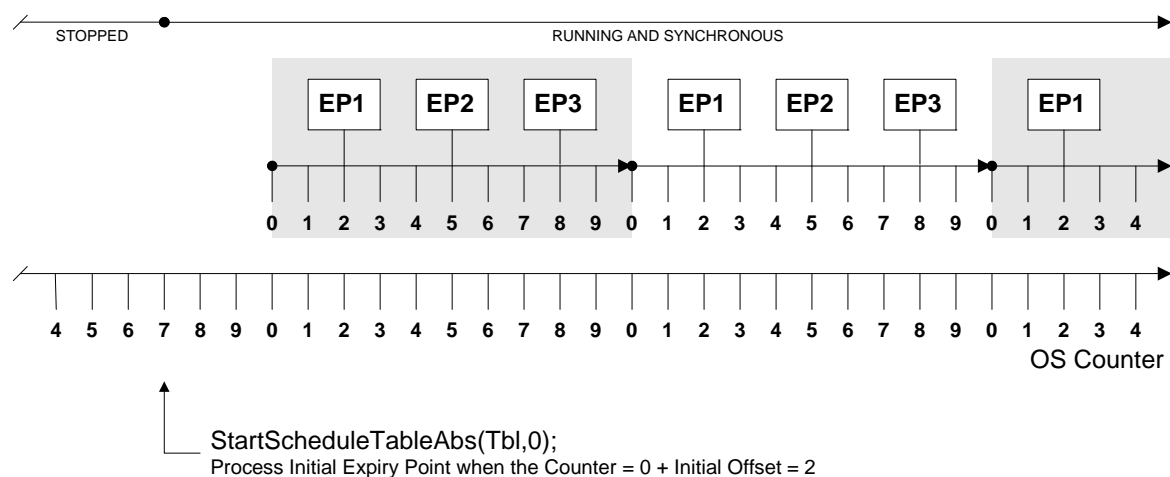


Figure 7.6: Example for implicit synchronized schedule table

7.4.2.2 Explicit Synchronization

An explicitly synchronized schedule table requires additional support from the Operating System module. The schedule table is driven by an Operating System module's counter as normal (termed the "drive counter") but processing needs to be synchronized with a different counter (termed the "synchronization counter") which is not an Operating System module's counter object.

The following constraints must be enforced between the schedule table, the Operating System module's counter and the synchronization counter:

Constraint1:

[OS431] [A schedule table that is explicitly synchronized shall have a duration no greater than modulus of the drive counter.] ()

Constraint2:

[OS462] [A schedule table that is explicitly synchronized shall have a duration equal to the modulus of the synchronization counter.] ()

Constraint3:

[OS463] [The synchronization counter shall have the same resolution as the drive counter associated with the schedule table. This means that a tick on the schedule table has the same duration as a tick on the synchronization counter.] ()

Note that it is in the responsibility of the Operating System module user to verify that Constraints 2 and 3 are satisfied by their system.

The function of explicit synchronization is for the Operating System module to keep processing each expiry point at absolute value of the synchronization counter equal to the expiry point's offset. This means that explicit synchronization always assumes that the notional zero of the schedule table has to be synchronized with absolute value zero on the synchronization counter.

To achieve this, the Operating System module must be told the value of the synchronization counter by the user. As the modulus of the synchronization counter and the schedule table are identical, the Operating System module can use this information to calculate drift. The Operating System module then automatically adjusts the delay between specially configured expiry points, retarding them or advancing them as appropriate, to ensure that synchronization is maintained.

7.4.2.2.1 Startup

There are two options for starting an explicitly synchronized schedule table:

1. Asynchronous start: Start the schedule table at an arbitrary value of the synchronization counter.
2. Synchronous start: Start the schedule table at absolute value zero of the synchronization counter only after a synchronization count has been provided. This may mean waiting for first synchronization indefinitely.

Asynchronous start is provided by the existing absolute and relative schedule table start services. Both of these services set the point at which the initial expiry point is processed with respect to the driver counter not the synchronization counter. This allows the schedule table to start running before the value of the synchronization counter is known.

Synchronous start requires an additional service that starts the schedule table only after the Operating System module is told the value of the synchronization counter.

The Operating System module provides the service `StartScheduleTableSynchron()` (see [OS201](#)) to start an explicitly synchronized

schedule table synchronously. The Initial Expiry Point will be processed after $(\text{Duration} - \text{Value}) + \text{Initial Offset}$ ticks of the driver counter have elapsed where Value is the absolute value of the synchronization counter provided to the schedule table.

[OS435] [If an explicitly synchronized schedule table was started synchronously, then the Operating System module shall guarantee that it has state “waiting” when the call of service `StartScheduleTableSynchron()` returns.] ()

7.4.2.2.2 Providing a Synchronization Count

The Operating System module must be told the value of the synchronization counter. Since the schedule table duration is equal to the modulus of the synchronization counter, the Operating System module can use this to determine the drift between the current count value on the schedule table time and the synchronization count and decide whether (or not) any action to achieve synchronization is required.

The Operating System module provides the service `SyncScheduleTable()` (see [OS199](#)) to provide the schedule table with a synchronization count and start synchronization.

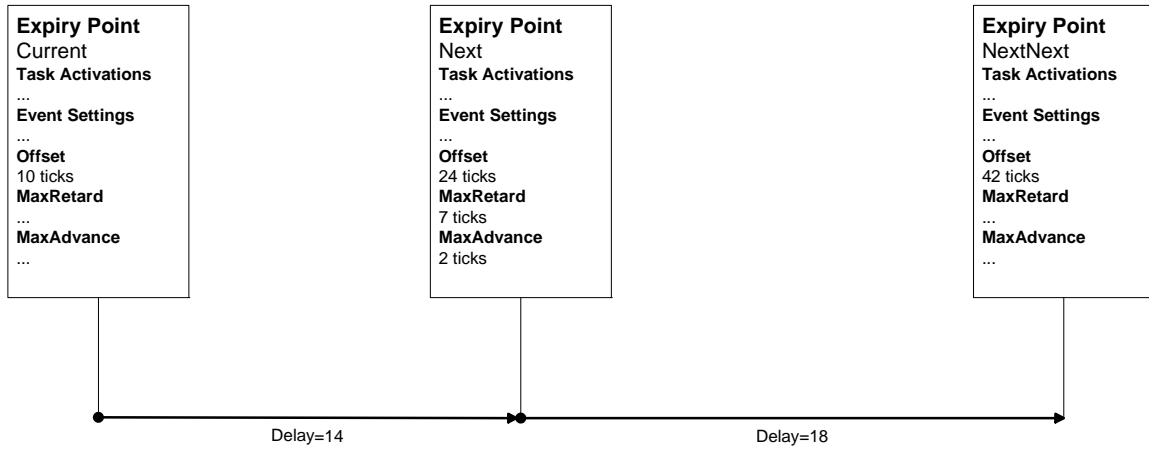
7.4.2.2.3 Specifying Synchronization Bounds

A schedule table defaults to denying adjustment at all expiry points. Adjustment is allowed only when explicitly configured. The range of adjustment that the Operating System module can make at an adjustable expiry point is controlled by specifying:

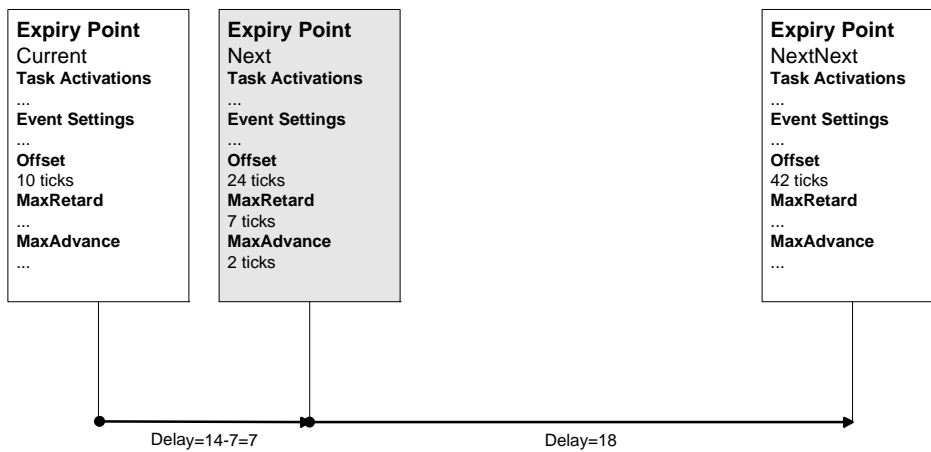
- `OsScheduleTableMaxShorten` : the maximum value that can be subtracted from the expiry offset
- `OsScheduleTableMaxLengthen`: the maximum value that can be added to the expiry point offset

The following figure illustrates the behaviour depending on `OsScheduleTableMaxShorten` and `OsScheduleTableMaxLengthen`:

Expected Delays



Maximum Retardation



Maximum Advance

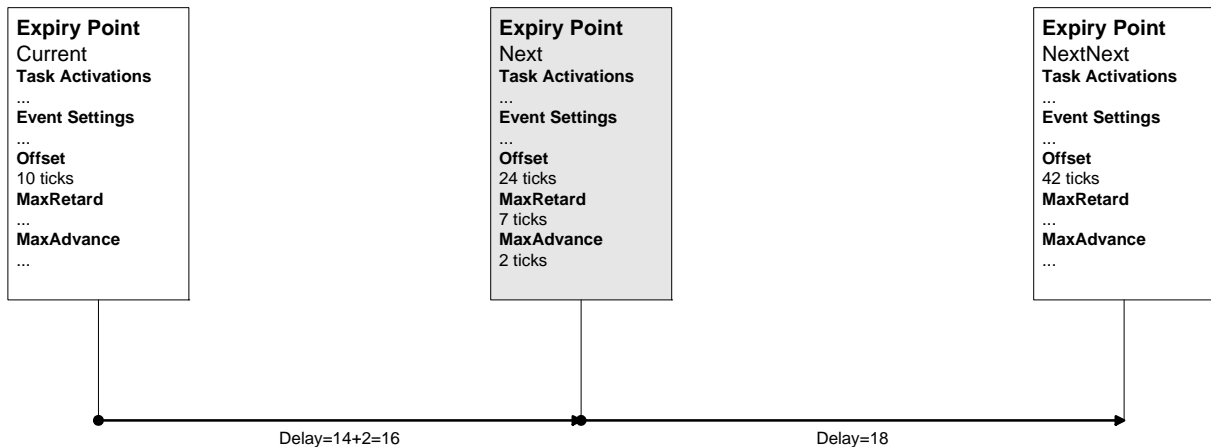


Figure 7.7: Adjustment of Expiry Points

So called “hard” and “smooth” synchronization from OSEKtime [16] are supported by this single unified concept in AUTOSAR OS. “Smooth” synchronization may be emulated by setting the small adjustment values on the final expiry point. “Hard” synchronization may be emulated by setting large adjustment values on the final expiry point.

[OS415] [An expiry point shall permit the configuration of a `OsScheduleTableMaxShorten` that defines the maximum number of ticks that can be subtracted from expiry point offset.] ()

[OS416] [An expiry point shall permit the configuration of a `OsScheduleTableMaxLengthen` that defines the maximum number of ticks that can be added to expiry point offset.] ()

When performing synchronization it is important that the expiry points on the schedule table are processed according to the total ordering defined by their offsets. This means that the range of permitted values for `OsScheduleTableMaxShorten` and `OsScheduleTableMaxLengthen` must ensure that the next expiry point is not retarded into the past or advanced beyond more than one iteration of the schedule table.

[OS436] [The value of $(\text{Offset} - \text{OsScheduleTableMaxShorten})$ of an expiry point shall be greater than $(\text{Offset} + \text{OsCounterMinCycle})$ of the previous expiry point.] ()

[OS559] [The value of `OsScheduleTableMaxLengthen` shall be smaller than the duration of the schedule table.] ()

[OS437] [The value of $(\text{OsScheduleTableMaxLengthen} + \text{delay_from_previous_EP})$ of an expiry point shall be less than the `OsCounterMaxAllowedValue` of the underlying counter.] ()

Explicitly synchronized schedule tables allow the tolerance of some drift between the schedule table value and the synchronization counter value. This tolerance can be zero, indicating that the schedule table is not considered synchronized unless the values are identical..

[OS438] [A schedule table shall define a precision bound with a value in the range 0 to duration.] ()

7.4.2.3 Performing Synchronization

The Operating System module uses the synchronization count to support (re-)synchronization of a schedule table at each expiry point by calculating an adjustment to the delay to the next expiry point. This provides faster re-synchronization of the schedule table than doing the action on the final expiry point.

[OS206] [When a new synchronization count is provided, the Operating System module shall calculate the current deviation between the explicitly synchronized scheduled table and the synchronization count.] (BSW11002)

It is meaningless to try and synchronise an explicitly synchronized schedule table before a synchronization count is provided.

[OS417] [The Operating System module shall start to synchronise an explicitly synchronized schedule table after a synchronization count is provided AND shall continue to adjust expiry points until synchronized.] ()

[OS418] [The Operating System module shall set the state of an explicitly synchronized schedule table to “running and synchronous” if the deviation is less than or equal to the configured `OsScheduleTblExplicitPrecision` threshold.] ()

[OS419] [The Operating System module shall set the state of an explicitly synchronized schedule table to “running” if the deviation is greater than the configured `OsScheduleTblExplicitPrecision` threshold.] ()

[OS420] [IF the deviation is non-zero AND the next expiry point is adjustable AND the table is behind the sync counter (`TableTicksAheadOfSyncCounter` \leq `TableTicksBehindOfSyncCounter`) THEN the OS shall set the next EP to expire delay - $\min(\text{MaxShorten}, \text{Deviation})$ ticks from the current expiry.] ()

[OS421] [IF the deviation is non-zero AND the next expiry point is adjustable AND the table is ahead of the sync counter (`TableTicksAheadOfSyncCounter` $>$ `TableTicksBehindOfSyncCounter`) THEN the OS shall set the next EP to expire delay + $\min(\text{MaxLengthen}, \text{Deviation})$ ticks from the current expiry.] ()

Figure 7.8: shows explicit synchronization of a schedule table. It assumes the following:

- EP1-3 have `OsScheduleTableMaxLengthen=2`
- EP1-3 have `OsScheduleTableMaxShorten =1`

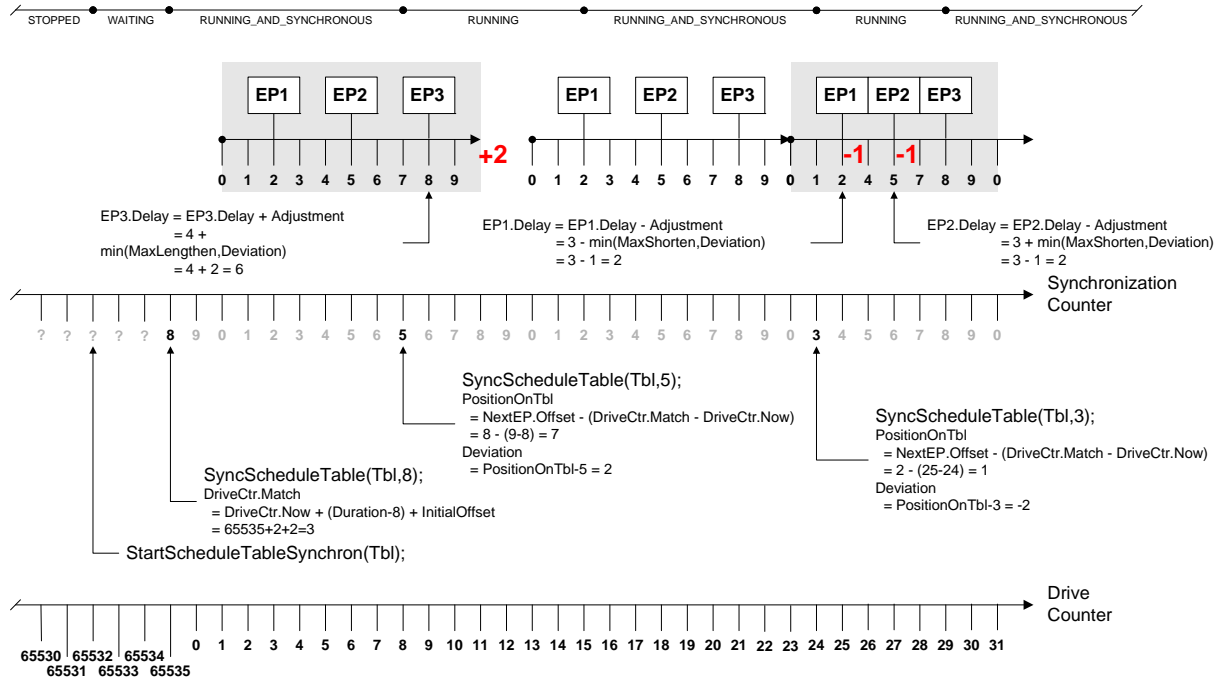


Figure 7.8: Explicit Schedule Table Synchronization

The Operating System module provides the service `SetScheduleTableAsync()` (see [OS422](#)) to cancel synchronization being performed at adjustable expiry points on a schedule table.

The Operating System module provides the service `GetScheduleTableStatus()` (see [OS227](#)) to query the state of a schedule table also with respect to synchronization.

7.5 Stack Monitoring Facilities

7.5.1 Background & Rationale

On processors that do not provide any memory protection hardware it may still be necessary to provide a “best effort with available resources” scheme for detectable classes of memory faults. Stack monitoring will identify where a task or ISR has exceeded a specified stack usage at context switch time. This may mean that there is considerable time between the system being in error and that fault being detected. Similarly, the error may have been cleared at the point the fault is notified (the stack may be less than the specified size when the context switch occurs).

It is not usually sufficient to simply monitor the entire stack space for the system because it is not necessarily the Task/ISR that was executing that used more than stack space than required – it could be a lower priority object that was pre-empted.

Significant debugging time can be saved by letting the Operating System correctly identify the Task/Category 2 ISR in error.

Note that for systems using a MPU and scalability class 3 or 4 a stack overflow may cause a memory exception before the stack monitoring is able to detect the fault.

7.5.2 Requirements

[OS067] [The Operating System module shall provide a stack monitoring which detects possible stack faults of Task(s)/Category 2 ISR(s).] (BSW11003)

[OS068] [If a stack fault is detected by stack monitoring AND no `ProtectionHook()` is configured, the Operating System module shall call the `ShutdownOS()` service with the status `E_OS_STACKFAULT`.] (BSW11003, BSW11013)

[OS396] [If a stack fault is detected by stack monitoring AND a `ProtectionHook()` is configured the Operating System module shall call the `ProtectionHook()` with the status `E_OS_STACKFAULT`.] ()

7.6 OS-Application

7.6.1 Background & Rationale

An AUTOSAR OS must be capable of supporting a collection of Operating System objects (Tasks, ISRs, Alarms, Schedule tables, Counters) that form a cohesive functional unit. This collection of objects is termed an *OS-Application*.

The Operating System module is responsible for scheduling the available processing resource between the OS-Applications that share the processor. If OS-Application(s) are used, all Tasks, ISRs, Counters, Alarms and Schedule tables must belong to an OS-Application. All objects which belong to the same OS-Application have access to each other. The right to access objects from other OS-Applications may be granted during configuration. An event is accessible if the task for which the event can be set is accessible. Access means that these Operating System objects are allowed as parameters to API services.

There are two classes of OS-Application:

- (1) Trusted OS-Applications are allowed to run with monitoring or protection features disabled at runtime. They may have unrestricted access to memory, the Operating System module's API, and need not have their timing behaviour enforced at runtime. They are allowed to run in privileged mode when supported by the processor.

(2) Non-Trusted OS-Applications are not allowed to run with monitoring or protection features disabled at runtime. They have restricted access to memory, restricted access to the Operating System module's API and have their timing behaviour enforced at runtime. They are not allowed to run in privileged mode when supported by the processor.

It is assumed that the Operating System module itself is trusted.

There are services offered by the AUTOSAR OS which give the caller information about the access rights and the membership of objects. These services are intended to be used in case of an inter-OS-Application call for checking access rights and arguments.

Note that Resource objects do not belong to any OS-Application, but access to them must be explicitly granted. (The same principle applies to spinlocks in Multi-Core systems)

The running OS-Application is defined as the OS-Application to which the currently running Task or ISR belongs. In case of a hook routine the Task or ISR which caused the call of the hook routine defines the running OS-Application.

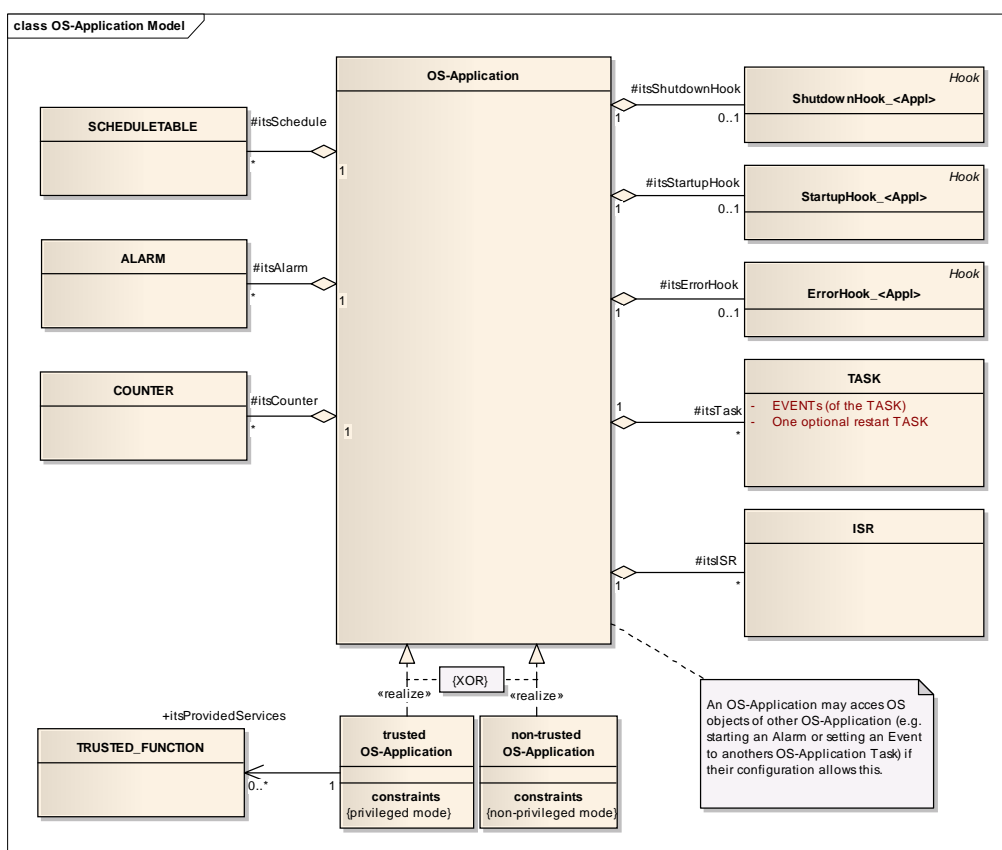


Figure 7.9: UML-model of OS-Application

OS-Applications have a state which defines the scope of accessibility of its Operating System objects from other OS-Applications. Each OS-Application is always in one of the following states:

- Active and accessible (APPLICATION_ACCESSIBLE): Operating System objects may be accessed from other OS-Applications. This is the default state at startup.
- Currently in restart phase (APPLICATION_RESTARTING). Operating System objects can not be accessed from other OS-Applications. State is valid until the OS-Application calls AllowAccess().
- Terminated and not accessible (APPLICATION_TERMINATED): Operating System objects can not be accessed from other OS-Applications. State will not change.

The following figure shows the states and the possible transitions:

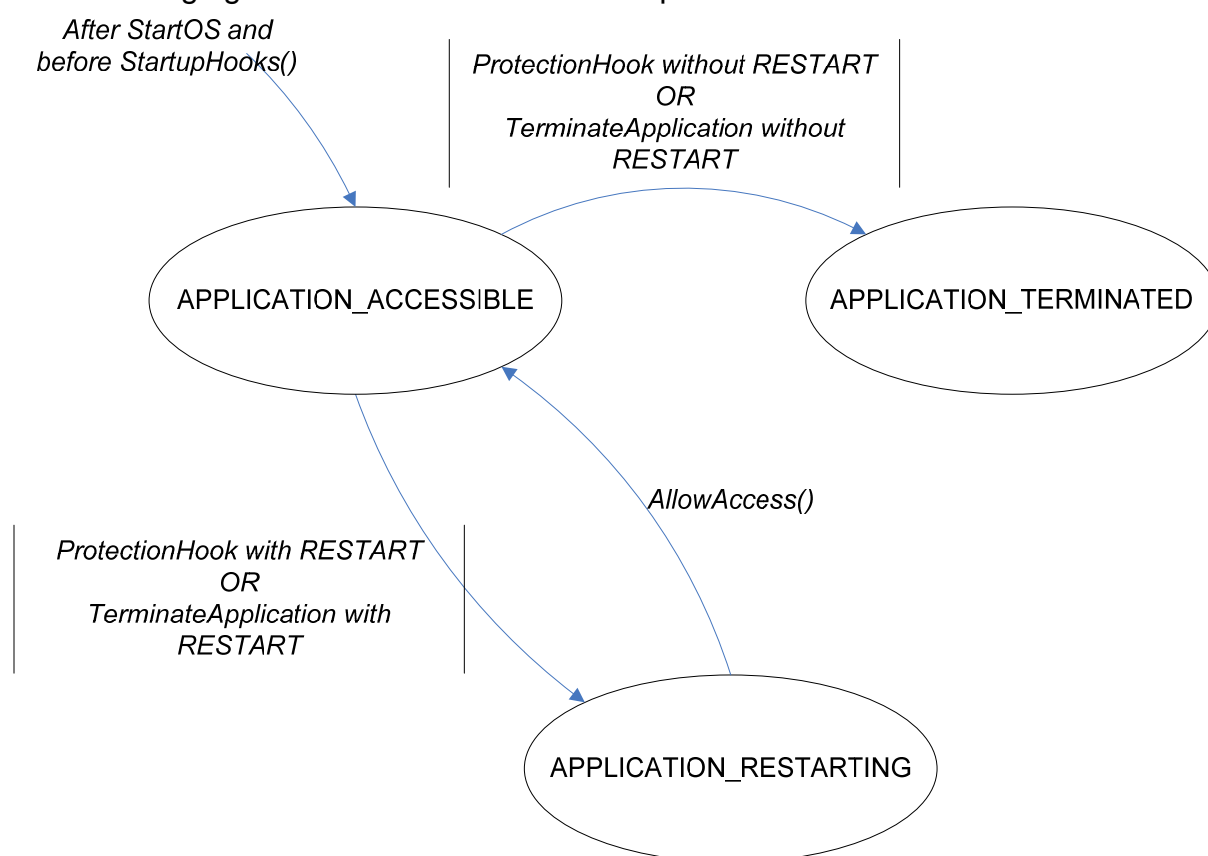


Figure 7.13: States of OS-Applications

7.6.2 Requirements

[OS445] [The Operating System module shall support OS-Applications which are a configurable selection of Trusted Functions, Tasks, ISRs, Alarms, Schedule tables, Counters, hooks (for startup, error and shutdown).] ()

[OS446] [The Operating System module shall support the notion of trusted and non-trusted OS-Applications.] ()

[OS464] [Trusted OS-Applications may offer services (“trusted services”) to other (even non-trusted) OS-Applications.] ()

The Operating System module provides the service `GetApplicationID()` (see [OS016](#)) to determine the currently running OS-Application (a unique identifier shall be allocated to each application).

The Operating System module provides the service `CheckObjectOwnership()` (see [OS017](#)) to determine to which OS-Application a given Task, ISR, Counter, Alarm or Schedule Table belongs.

The Operating System module provides the service `CheckObjectAccess()` (see [OS256](#)) to determine which OS-Applications are allowed to use the IDs of a Task, ISR, Resource, Counter, Alarm or Schedule Table in API calls.

The Operating System module provides the service `TerminateApplication()` (see [OS258](#)) to terminate the OS-Application to which the calling Task/Category 2 ISR/application specific error hook belongs. (This is an OS-Application level variant of the `TerminateTask()` service)

The Operating System provides the service `TerminateApplication()` (see [OS258](#)) to terminate another OS-Application AND calls to this service shall be ignored if the caller does not belong to a trusted OS-Application.

[OS447] [If the Operating System module terminates an OS-Application, then it shall:

- terminate all running, ready and waiting Tasks/ISRs of the OS-Application AND
- disable all interrupts of the OS-Application AND
- stop all active alarms of the OS-Applications AND
- stop all schedule tables of the OS-Application.] ()

[OS448] [The Operating System module shall prevent access of OS-Applications, trusted or non-trusted, to objects not belonging to this OS-Application, except access rights for such objects are explicitly granted by configuration.] ()

The Operating System provides the service `GetApplicationState()` (see [OS499](#)) to request the current state of an OS-Application.

[OS500] [The Operating System module shall set the state of all OS-Applications after the call of `StartOS()` and before any `StartupHook` is called to `APPLICATION_ACCESSIBLE`.] ()

The Operating System module provides the service `AllowAccess()` (see [OS501](#)) to set the own state of an OS-Application from `APPLICATION_RESTARTING` to `APPLICATION_ACCESSIBLE`.

[OS502] [If an OS-Application is terminated (e.g. through a service call or via protection hook) and no restart is requested, then the Operating System module shall set the state of this OS-Application to `APPLICATION_TERMINATED`.] ()

[OS503] [If an OS-Application is terminated (e.g. through a service call or via protection hook) and a restart is requested, then the Operating System module shall set the state of this OS-Application to `APPLICATION_RESTARTING`.] ()

[OS504] [The Operating System module shall deny access to Operating System objects from other OS-Applications to an OS-Application which is not in state `APPLICATION_ACCESSIBLE`.] ()

[OS509] [If a service call is made on an Operating System object that is owned by another OS-Application without state `APPLICATION_ACCESSIBLE`, then the Operating System module shall return `E_OS_ACCESS`.] ()

An example for [OS509](#) is a call to `ActivateTask()` for a task in an OS-Application that is restarting.

7.7 Protection Facilities

Protection is only possible for Operating System managed objects. This means that:

- It is not possible to provide protection during runtime of Category 1 ISRs, because the operating system is not aware of any Category 1 ISRs being invoked. Therefore, if any protection is required, Category 1 ISRs have to be avoided. If Category 1 interrupts AND OS-Applications are used together then all Category 1 ISR must belong to a trusted OS-Application.
- It is not possible to provide protection between functions called from the body of the same Task/Category 2 ISR.

7.7.1 Memory Protection

7.7.1.1 Background & Rationale

Memory protection will only be possible on processors that provide hardware support for memory protection.

The memory protection scheme is based on the (data, code and stack) sections of the executable program.

Stack: An OS-Application comprises a number of Tasks and ISRs. The stack for these objects, by definition, belongs only to the owner object and there is therefore no need to share stack data between objects, even if those objects belong to the same OS-Application.

Memory protection for the stacks of Tasks and ISRs is useful mainly for two reasons:

- (1) Provide a more immediate detection of stack overflow and underflow for the Task or ISR than can be achieved with stack monitoring
- (2) Provide protection between constituent parts of and OS-Application, for example to satisfy some safety constraints.

Data: OS-Applications can have private data sections and Tasks/ISRs can have private data sections. OS-Application's private data sections are shared by all Tasks/ISRs belonging to that OS-Application.

Code: Code sections are either private to an OS-Application or can be shared between all OS-Applications (to use shared libraries). In the case where code protection is not used, executing incorrect code will eventually result in a memory, timing or service violation.

7.7.1.2 Requirements

Data Sections and Stack

[OS198] [The Operating System module shall prevent write access to its own data sections and its own stack from non-trusted OS-Applications.] ()

Private data of an OS-Application

[OS026] [The Operating System module may prevent read access to an OS-Application's data section attempted by other non-trusted OS-Applications.] (BSW11000)

[OS086] [The Operating System module shall permit an OS-Application read and write access to that OS-Application's own private data sections.] (BSW11006)

[OS207] [The Operating System module shall prevent write access to the OS-Application's private data sections from other non-trusted OS-Applications.] (BSW11005)

Private Stack of Task/ISR

[OS196] [The Operating System module shall permit a Task/Category 2 ISR read and write access to that Task's/Category 2 ISR's own private stack.] (BSW11006)

[OS208] [The Operating System module may prevent write access to the private stack of Tasks/Category 2 ISRs of a non-trusted application from all other Tasks/ISRs in the same OS-Application.] (BSW11005)

[OS355] [The Operating System module shall prevent write access to all private stacks of Tasks/Category 2 ISRs of an OS-Application from other non-trusted OS-Applications.] ()

Private data of a Task/ISR

[OS087] [The Operating System module shall permit a Task/Category 2 ISR read and write access to that Task's/Category 2 ISR's own private data sections.] (BSW11006)

[OS195] [The Operating System module may prevent write access to the private data sections of a Task/Category 2 ISR of a non-trusted application from all other Tasks/ISRs in the same OS-Application.] (BSW11005)

[OS356] [The Operating System module shall prevent write access to all private data sections of a Task/Category 2 ISR of an OS-Application from other non-trusted OS-Applications.] ()

Code Sections

[OS027] [The Operating System module may provide an OS-Application the ability to protect its code sections against executing by non-trusted OS-Applications.] ()

[OS081] [The Operating System module shall provide the ability to provide shared library code in sections that are executable by all OS-Applications.] (BSW11007)

Peripherals

[OS209] [The Operating System module shall permit trusted OS-Applications read and write access to peripherals.] ()

[OS083] [The Operating System module shall allow non-trusted OS-Applications to write to their assigned peripherals only (incl. reads that have the side effect of writing to a memory location).] ()

Memory Access Violation

[OS044] [If a memory access violation is detected, the Operating System module shall call the Protection Hook with status code `E_OS_PROTECTION_MEMORY`.] (BSW11013)

7.7.2 Timing Protection

7.7.2.1 Background & Rationale

A timing fault in a real-time system occurs when a task or interrupt misses its deadline at runtime.

AUTOSAR OS does not offer deadline monitoring for timing protection. Deadline monitoring is insufficient to correctly identify the Task/ISR causing a timing fault in an AUTOSAR system. When a deadline is violated this may be due to a timing fault introduced by an unrelated Task/ISR that interferes/blocks for too long. The fault in this case lies with the unrelated Task/ISR and this will propagate through the system until a Task/ISR misses its deadline. The Task/ISR that misses a deadline is therefore not necessarily the Task/ISR that has failed at runtime, it is simply the earliest point that a timing fault is detected.

If action is taken based on a missed deadline identified with deadline monitoring this would potentially use false evidence of error to terminate a correct OS-Application in favour of allowing an incorrect OS-Application to continue running. The problem is best illustrated by example. Consider a system with the following configuration:

TaskID	Priority	Execution Time	Deadline (=Period)
A	High	1	5
B	Medium	3	10
C	Low	5	15

Assuming that all tasks are ready to run at time zero, the following execution trace would be expected and all tasks would meet their respective deadlines.




Figure 7.10: Example execution trace

Now consider the case when tasks A and B behave incorrectly. The figure below shows both task A and task B executing for longer than specified and task B arriving 2 ticks earlier than specified. Both tasks A and B meet their deadlines. Task C however, behaves correctly but it fails to meet its deadline because of the incorrect execution of Tasks A and B. This is fault propagation – a fault in an unrelated part of the system is causing a correctly functioning part of the system to fail.

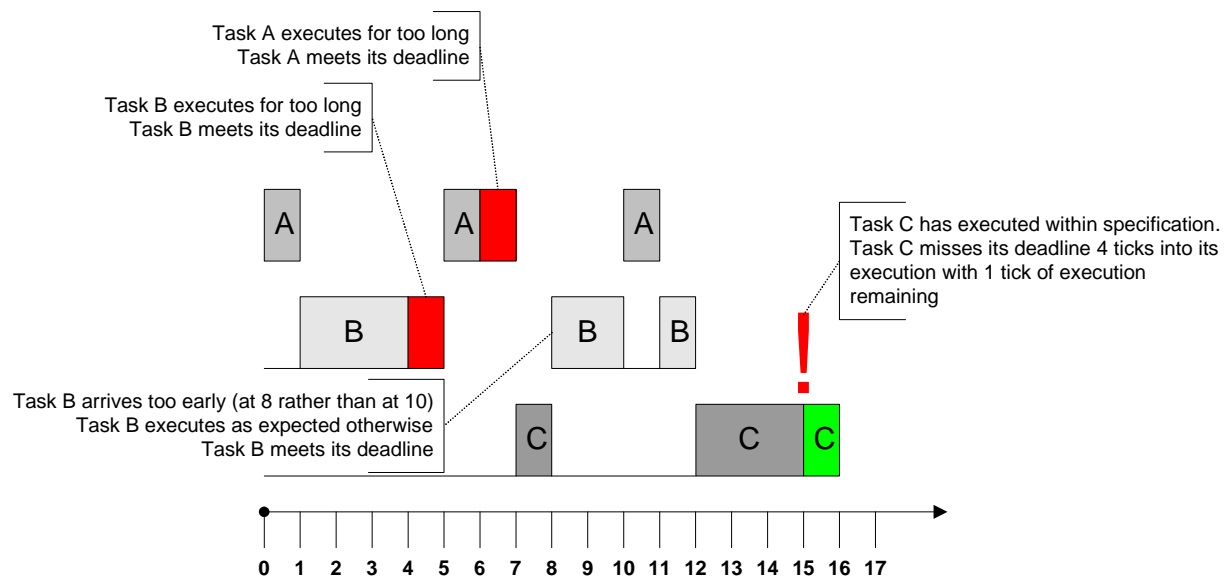


Figure 7.11: Insufficiency of Deadline Monitoring

Whether a task or ISR meets its deadline in a fixed priority preemptive operating system like AUTOSAR OS is determined by the following factors:

- (1) the execution time of Task/ISRs in the system
- (2) the blocking time that Task/ISRs suffers from lower priority Tasks/ISRs locking shared resources or disabling interrupts
- (3) the interarrival rate of Task/ISRs in the system

For safe and accurate timing protection it is necessary for the operating system to control these factors at runtime to ensure that Tasks/ISRs can meet their respective deadlines.

AUTOSAR OS prevents timing errors from (1) by using *execution time protection* to guarantee a statically configured upper bound, called the Execution Budget, on the execution time of:

- Tasks
- Category 2 ISRs

AUTOSAR OS prevents timing errors from (2) by using *locking time protection* to guarantee a statically configured upper bound, called the Lock Budget, on the time that:

- Resources are held by Tasks/Category 2 ISRs
- OS interrupts are suspended by Tasks/Category 2 ISRs
- ALL interrupts are suspended/disabled by Tasks/Category 2 ISRs

AUTOSAR OS prevents timing errors from (3) by using *inter-arrival time protection* to guarantee a statically configured lower bound, called the Time Frame, on the time between:

- A task being permitted to transition into the `READY` state due to:
 - Activation (the transition from the `SUSPENDED` to the `READY` state)
 - Release (the transition from the `WAITING` to the `READY` state)
- A Category 2 ISR arriving
An arrival occurs when the Category 2 ISR is recognized by the OS

Inter-arrival time protection for basic tasks controls the time between successive activations, irrespective of whether activations are queued or not. In the case of queued activations, activating a basic task which is in the `READY` or `RUNNING` state is a new activation because it represents the activation of a new instance of the task. Inter-arrival time protection therefore interacts with queued activation to control the rate at which the queue is filled.

Inter-arrival time protection for extended tasks controls the time between successive activations *and* releases. When a task is in the `WAITING` state and multiple events are set with a single call to `SetEvent()` this represents a single release. When a task waits for one or more events which are already set this represents a notional Wait/Release/Start transition and therefore is considered as a new release.

The following figure shows how execution time protection and inter-arrival time protection interact with the task state transition model for AUTOSAR OS.

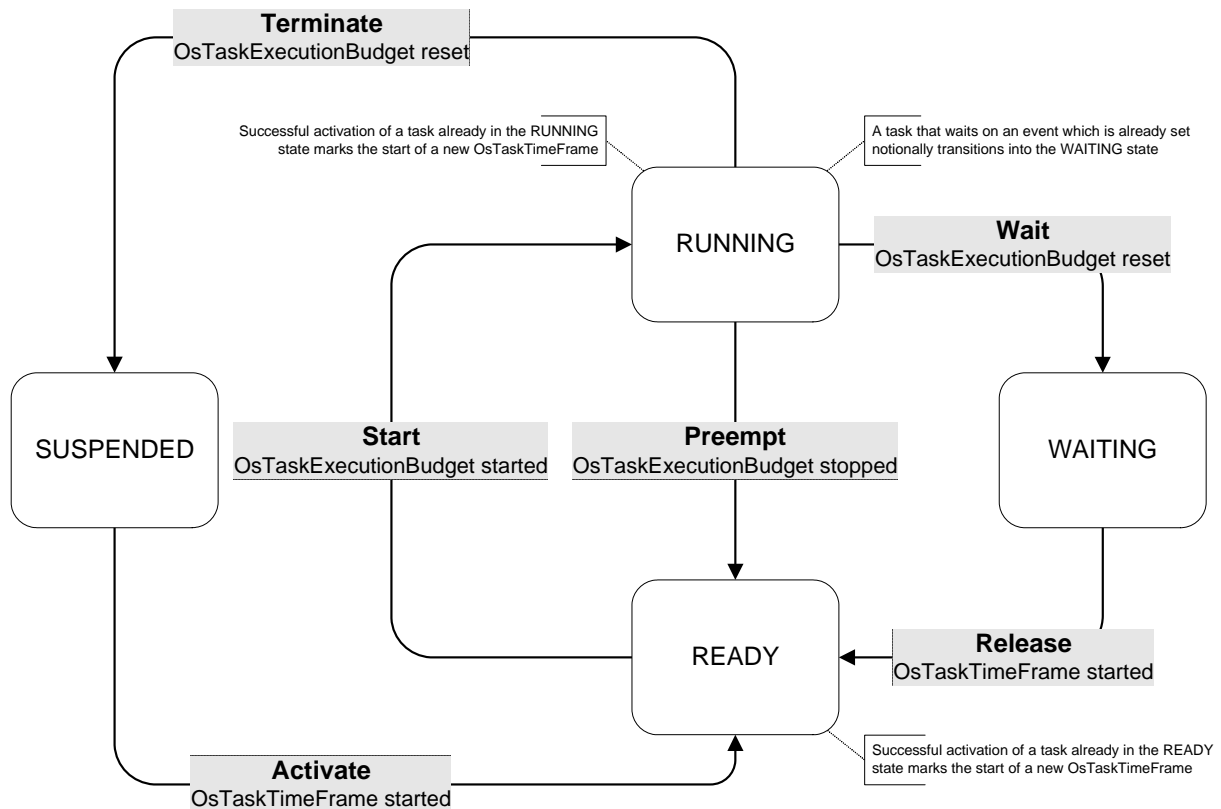


Figure 7.12: Time protection interaction with the task state transition model

Notes:

1. Inter-arrival time enforcement on Category 2 ISRs can be used to protect an ECU from a “babbling idiot” source of interrupts (e.g. a CAN controller taking an interrupt each time a frame is received from another ECU on the network) and provides the type of protection given by the OSEKtime Interrupt re-enable schedule event [16].
2. Timing protection only applies to Tasks or Category 2 ISRs. There is no protection for Category 1 ISRs. If timing protection error occurs during a category 1 ISR, consistency of the Operating System module can not be guaranteed. Therefore we discourage timing protection in systems with category 1 interrupts.
3. Timing protection does not apply before the Operating System module is started.
4. In the case of trusted OS-Applications it is essential that all timing information is correct, otherwise the system may fail at run-time. For a non-trusted OS-Application, timing protection can be used to enforce timing boundaries between executable objects.

7.7.2.2 Requirements

[OS028] [In a non-trusted OS-Application, the Operating System module shall apply timing protection to every Task/Category 2 ISR of this non-trusted OS-Application.] (BSW11008)

[OS089] [In a trusted OS-Application, the Operating System module shall provide the ability to apply timing protection to Tasks/Category 2 ISRs of this OS-Application.] (BSW11008)

[OS397] [If no OS-Application is configured, the Operating System module shall be able to apply timing protection to Tasks/Category 2 ISRs.] ()

Timing Protection: Tasks

[OS064] [If a task's `OsTaskExecutionBudget` is reached then the Operating System module shall call the `ProtectionHook()` with `E_OS_PROTECTION_TIME`.] (BSW11008, BSW11013)

[OS473] [The Operating System module shall reset a task's `OsTaskExecutionBudget` on a transition to the `SUSPENDED` or `WAITING` states.] (BSW11008)

[OS465] [The Operating System module shall limit the inter-arrival time of tasks to one per `OsTaskTimeFrame`.] (BSW11008)

[OS469] [The Operating System module shall start an `OsTaskTimeFrame` when a task is activated successfully.] (BSW11008)

[OS472] [The Operating System module shall start an `OsTaskTimeFrame` when a task is released successfully.] (BSW11008)

[OS466] [If an attempt is made to activate a task before the end of an `OsTaskTimeFrame` then the Operating System module shall not perform the activation AND shall call the `ProtectionHook()` with `E_OS_PROTECTION_ARRIVAL`.] ()

[OS467] [If an attempt is made to release a task before the end of an `OsTaskTimeFrame` then the Operating System module shall not perform the release AND shall call the `ProtectionHook()` with `E_OS_PROTECTION_ARRIVAL` AND the event shall be set.] ()

Timing Protection: ISRs

[OS210] [If a Category 2 ISR's `OsIsrExecutionBudget` is reached then the Operating System module shall call the `ProtectionHook()` with `E_OS_PROTECTION_TIME`.] (BSW11013)

[OS474] [The Operating System module shall reset an ISR's `OslsrExecutionBudget` when the ISR returns control to the Operating System.] (BSW11008)

[OS470] [The Operating System module shall limit the inter-arrival time of Category 2 ISRs to one per `OslsrTimeFrame`.] (BSW11008)

[OS471] [The Operating System module shall measure the start of an `OslsrTimeFrame` from the point at which it recognises the interrupt (i.e. in the Operating System interrupt wrapper).] (BSW11008)

[OS048] [If Category 2 interrupt occurs before the end of the `OslsrTimeFrame` then the Operating System module shall not execute the user provided ISR AND shall call the `ProtectionHook()` with `E_OS_PROTECTION_ARRIVAL`.] (BSW11008)

Timing Protection: Resource Locking and Interrupt Disabling

[OS033] [If a Task/Category 2 ISR holds an OSEK Resource and exceeds the `Os[Task|Isr]ResourceLockBudget`, the Operating System module shall call the `ProtectionHook()` with `E_OS_PROTECTION_LOCKED`.] (BSW11008, BSW11013, BSW11014)

[OS037] [If a Task/Category 2 ISR disables interrupts (via `Suspend/DisableAll/OSInterrupts()`) and exceeds the configured `Os[Task|Isr][All|OS]InterruptLockBudget`, the Operating System module shall call the `ProtectionHook()` with `E_OS_PROTECTION_LOCKED`.] (BSW11008, BSW11013, BSW11014)

7.7.2.3 Implementation Notes

Execution time enforcement requires hardware support, e.g. a timing enforcement interrupt. If an interrupt is used to implement the time enforcement, the priority of this interrupt has to be high enough to “interrupt” the supervised tasks or ISRs.

Depending on the real hardware support this could mean that `DisableAllInterrupts` and `SuspendAllInterrupts` disable not all interrupts (e.g. all interrupts except of the interrupt used for timing protection) or that the usage of Category 1 ISRs – which bypass the Operating System (and also the timing protection) – is limited somehow.

The implementation has to document such implementation specific behaviour (e.g. the limitations when timing protection is used).

7.7.3 Service Protection

Background & Rationale

As OS-Applications can interact with the Operating System module through services, it is essential that the service calls will not corrupt the Operating System module itself. Service Protection guards against such corruption at runtime.

There are a number of cases to consider with Service Protection: An OS-Application makes an API call

- (1) with an invalid handle or out of range value.
- (2) in the wrong context, e.g. calling `ActivateTask()` in the `StartupHook()`.
- (3) or fails to make an API call that results in the OSEK OS being left in an undefined state, e.g. it terminates without a `ReleaseResource()` call
- (4) that impacts on the behaviour of every other OS-Application in the system, e.g. `ShutdownOS()`
- (5) to manipulate Operating System objects that belong to another OS-Application (to which it does not have the necessary permissions), e.g. an OS-Application tries to execute `ActivateTask()` on a task it does not own.

The OSEK OS already provides some service protection through the status codes returned from service calls and this will provide the basis for service protection. This means that service protection will only apply for the extended status of OSEK OS.

However, OSEK OS does not cover all the cases outlined above. The following sections describe – besides the mandatory extended status – the additional protection requirements to be applied in each of these cases.

7.7.3.1 Invalid Object Parameter or Out of Range Value

7.7.3.1.1 Background & Rationale

The current OSEK OS' service calls already return `E_OS_ID` on invalid objects (i.e. objects not defined in the OIL file) and `E_OS_VALUE` for out of range values (e.g. setting an alarm cycle time less than `OsCounterMinCycle`).

7.7.3.1.2 Requirements

[OS051] [If an invalid address (address is not writable by this OS-Application) is passed as an out-parameter to an Operating System service, the Operating System module shall return the status code `E_OS_ILLEGAL_ADDRESS`.] (BSW11009, BSW11013)

7.7.3.2 Service Calls Made from Wrong Context

7.7.3.2.1 Background & Rationale

The current OSEK OS defines the valid calling context for service calls ([15], Fig. 12-1), however protects against only a small set of these invalid calls, e.g. calling `TerminateTask()` from a Category 2 ISR.

Service	Task	Cat1 ISR	Cat2 ISR	Error Hook	PreTask Hook	PostTask Hook	Startup Hook	Shutdown Hook	Alarm Callback	Protection Hook
ActivateTask	✓		✓							
TerminateTask	✓		C							
ChainTask	✓		C							
Schedule	✓		C							
GetTaskID	✓		✓	✓	✓	✓				✓
GetTaskState	✓		✓	✓	✓	✓				
DisableAllInterrupts	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
EnableAllInterrupts	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
SuspendAllInterrupts	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
ResumeAllInterrupts	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
SuspendOSInterrupts	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
ResumeOSInterrupts	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
GetResource	✓		✓							
ReleaseResource	✓		✓							
SetEvent	✓		✓							
ClearEvent	✓		C							
GetEvent	✓		✓	✓	✓	✓				
WaitEvent	✓		C							
GetAlarmBase	✓		✓	✓	✓	✓				
GetAlarm	✓		✓	✓	✓	✓				
SetRelAlarm	✓		✓							
SetAbsAlarm	✓		✓							
CancelAlarm	✓		✓							
GetActiveApplicationMode	✓		✓	✓	✓	✓	✓	✓		
StartOS										
ShutdownOS	✓		✓	✓			✓			
GetApplicationID	✓		✓	✓	✓	✓	✓	✓		✓
GetISRID	✓		✓	✓						✓
CallTrustedFunction	✓		✓							
CheckISRMemoryAccess	✓		✓	✓						✓
CheckTaskMemoryAccess	✓		✓	✓						✓
CheckObjectAccess	✓		✓	✓						✓
CheckObjectOwnership	✓		✓	✓						✓
StartScheduleTableRel	✓		✓							

Service	Task	Cat1 ISR	Cat2 ISR	Error Hook	PreTask Hook	PostTask Hook	Startup Hook	Shutdown Hook	Alarm Callback	Protection Hook
StartScheduleTableAbs	✓		✓							
StopScheduleTable	✓		✓							
NextScheduleTable	✓		✓							
StartScheduleTableSynchron	✓		✓							
SyncScheduleTable	✓		✓							
GetScheduleTableStatus	✓		✓							
SetScheduleTableAsync	✓		✓							
IncrementCounter	✓		✓							
GetCounterValue	✓		✓							
GetElapsedValue	✓		✓							
TerminateApplication	✓		✓	✓ ²						
AllowAccess	✓		✓							
GetApplicationState	✓		✓	✓	✓	✓	✓	✓		✓

Tab. 1: Allowed Calling Context for OS Service Calls

In the table above “C” indicates that validity is only “Checked in Extended status by E_OS_CALLEVEL” (see [12], section 13.1).

7.7.3.2.2 Requirements

[OS088] [If an OS-Application makes a service call from the wrong context AND is currently not inside a Category 1 ISR the Operating System module shall not perform the requested action (the service call shall have no effect), and return E_OS_CALLEVEL (see [12], section 13.1) or the “invalid value” of the service.] (BSW11009, BSW11013)

7.7.3.3 Services with Undefined Behaviour

7.7.3.3.1 Background & Rationale

There are a number of situations where the behaviour of OSEK OS is undefined in extended status. This is unacceptable when protection is required as it would allow the Operating System module to be corrupted through its own service calls. The implementation of service protection for the Operating System module must therefore describe and implement a behaviour that does not jeopardise the integrity of the system or of any OS-Application which did not cause the specific error.

² Only in case of self termination.
73 of 230

7.7.3.3.2 Requirements

Tasks ends without calling a `TerminateTask()` or `ChainTask()`

[OS052] [If a task returns from its entry function without making a `TerminateTask()` or `ChainTask()` call, the Operating System module shall terminate the task (and call the `PostTaskHook()` if configured).] (BSW11009)

[OS069] [If a task returns from its entry function without making a `TerminateTask()` or `ChainTask()` call AND the error hook is configured, the Operating System module shall call the `ErrorHook()` (this is done regardless of whether the task causes other errors, e.g. `E_OS_RESOURCE`) with status `E_OS_MISSINGEND` before the task leaves the `RUNNING` state.] (BSW11009)

[OS070] [If a task returns from the entry function without making a `TerminateTask()` or `ChainTask()` call and still holds OSEK Resources, the Operating System module shall release them.] (BSW11009, BSW11013)

[OS239] [If a task returns from the entry function without making a `TerminateTask()` or `ChainTask()` call and interrupts are still disabled, the Operating System module shall enable them.] ()

Category 2 ISR ends with locked interrupts or allocated resources

[OS368] [If a Category 2 ISR calls `DisableAllInterupts()` / `SuspendAllInterrupts()` / `SuspendOSInterrupts()` and ends (returns) without calling the corresponding `EnableAllInterrupts()` / `ResumeAllInterrupts()` / `ResumeOSInterrupts()`, the Operating System module shall perform the missing service and shall call the `ErrorHook()` (if configured) with the status `E_OS_DISABLEDINT`.] ()

[OS369] [If a Category 2 ISR calls `GetResource()` and ends (returns) without calling the corresponding `ReleaseResource()`, the Operating System module shall perform the `ReleaseResource()` call and shall call the `ErrorHook()` (if configured) with the status `E_OS_RESOURCE` (see [12], section 13.1).] ()

PostTaskHook called during ShutdownOS()

[OS071] [If the `PostTaskHook()` is configured, the Operating System module shall not call the hook if `ShutdownOS()` is called.] ()

Tasks/ISRs calls EnableAllInterrupts/ResumeAllInterrupts/ResumeOSInterrupts without a corresponding disable

[OS092] [If `EnableAllInterrupts()` / `ResumeAllInterrupts()` / `ResumeOSInterrupts()` are called and no corresponding `DisableAllInterupts()` / `SuspendAllInterrupts()` / `SuspendOSInterrupts()` was done before, the Operating System module shall not perform this Operating System service.] (BSW11009)

Tasks/ISRs calling OS services when DisableAllInterupts/SuspendAllInterrupts/SuspendOSInterrupts called

[OS093] [If interrupts are disabled/suspended by a Task/ISR/Hook and the Task/ISR/Hook calls any Operating System service (excluding the interrupt services) then the Operating System module shall ignore the service AND shall return `E_OS_DISABLEDINT` if the service returns a `StatusType` value.] (BSW11009, BSW11013)

7.7.3.4 Service Restrictions for Non-Trusted OS-Applications

7.7.3.4.1 Background & Rationale

The Operating System service calls available are restricted according to the calling context (see Section 7.7.3.2). In a protected system, additional constraints need to be placed to prevent non-trusted OS-Applications executing API calls that can have a global effect on the system. Each level of restriction is a proper subset of the previous level as shown in the figure below.

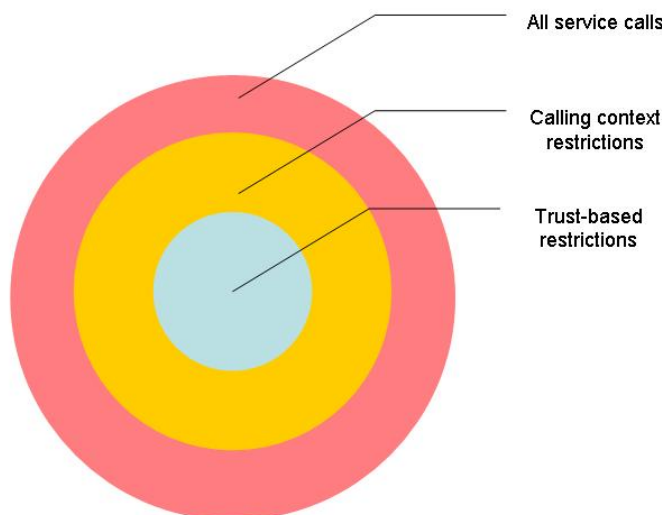


Figure 7.13: API Restrictions

There are two defined integrity levels:

1. Trusted
2. Non-Trusted

that correspond exactly with trusted and non-trusted OS-Applications.

7.7.3.4.2 Requirements

[OS054] [The Operating System module shall ignore calls to `ShutdownOS()` from non-trusted OS-Applications.] ()

7.7.3.5 Service Calls on Objects in Different OS-Applications

7.7.3.5.1 Background

Section 7.7.3.1 stated that `E_OS_ID` is returned by OSEK OS service calls when the object is invalid. Under the protection scheme a service call can be invalid because the caller does not have valid permissions for the object (a new meaning for multi-OS-Application systems).

This is a similar case to an object not being accessible in OSEK OS (for example, when a task tries to get a resource which exists in the system but has not been configured as used by the task).

7.7.3.5.2 Requirements

[OS056] [If an OS-object identifier is the parameter of an Operating System module's system service, and no sufficient access rights have been assigned to this OS-object at configuration time (Parameter `Os[...]AccessingApplication`) to the calling Task/Category 2 ISR, the Operating System module's system service shall return `E_OS_ACCESS`.] (BSW11001, BSW11010, BSW11013)

[OS449] [CheckTaskMemoryAccess and CheckIsrMemoryAccess check the memory access. Memory access checking is possible for all OS-Applications and from all OS-Applications and does not need granted rights.] ()

[OS449](#) is an exception to [OS056](#).

[OS450] [CheckObjectAccess checks the access rights for Operating System objects. Checking object access is possible for all OS-Applications and from all OS-Applications and does not need granted rights.] ()

[OS450](#) is an exception to [OS056](#).

7.7.4 Protecting the Hardware used by the OS

7.7.4.1 Background & Rationale

Where a processor supports privileged and non-privileged mode it is usually the case that certain registers, and the instructions to modify those registers, are inaccessible outside the privileged mode.

On such hardware, executing the Operating System module in privileged mode and Tasks/ISRs in non-privileged mode protects the registers fundamental to Operating System module operation from inadvertent corruption by the objects executing in non-privileged mode. The Operating System module's services will need to execute in privileged mode as they will need to modify the registers that are protected outside this mode.

The Operating System module can use the control registers of the MPU, timer unit(s), interrupt controller, etc. and therefore it is necessary to protect those registers against non-trusted OS-Applications.

7.7.4.2 Requirements

[OS058] [If supported by hardware, the Operating System module shall execute non-trusted OS-Applications in non-privileged mode.] ()

[OS096] [As far as supported by hardware, the Operating System module shall not allow non-trusted OS-Applications to access control registers managed by the Operating System module.] (BSW11011)

[OS245] [If an instruction exception occurs (e.g. division by zero) the Operating System module shall call the protection hook with `E_OS_PROTECTION_EXCEPTION`.] (BSW11011)

7.7.4.3 Implementation Notes

When the Operating System module is running non-trusted OS-Applications, the Operating System module's treatment of interrupt entry and hook routines must be carefully managed.

Interrupt handling: Where the MCU supports different modes (as discussed in this section) ISRs will require the Operating System module to do extra work in the `ISR()` wrapper. ISRs will typically be entered in privileged mode. If the handler is part of a non-trusted OS-Application then the `ISR()` wrapper must make sure that a switch to non-privileged mode occurs before the handler executes.

7.7.5 Providing »Trusted Functions«

7.7.5.1 Background & Rationale

An OS-Application can invoke a Trusted Function provided by (another) trusted OS-Application. That can require a switch from non-privileged to privileged mode. This is typically achieved by these operations:

- (1) Each trusted OS-Application may export services which are callable from other OS-Applications.
- (2) During configuration these trusted services must be configured to be called from a non-trusted OS-Application.
- (3) The call from the non-trusted OS-Application to the trusted service is using a mechanism (e.g. trap/software interrupt) provided by the Operating System. The service is passed as an identifier that is used to determine, in the trusted environment, if the service can be called.
- (4) The Operating System offers services to check if a memory region is write/read/execute accessible from an OS-Application. It also returns information if the memory region is part of the stack space.

The Operating System software specification does not provide support for »non-trusted services«.

7.7.5.2 Requirements

[OS451] [The Operating System module shall allow exporting services from trusted OS-Applications.] ()

The Operating System module provides the service `CallTrustedFunction()` (see [OS097](#)) to call a trusted function from a (trusted or non-trusted) OS-Application.

[OS100] [If `CallTrustedFunction()` is called and the called trusted function is not configured the Operating System module shall call the `ErrorHook` with `E_OS_SERVICEID`.] ()

The Operating System module provides the services `CheckISRMemoryAccess()` and `CheckTaskMemoryAccess()` (see [OS512](#) and [OS513](#)) for OS-Applications to check if a memory region is write/read/execute accessible from a Task/Category 2 ISR and also return information if the memory region is part of the stack space.

7.8 Protection Error Handling

7.8.1 Background & Rationale

The Operating System can detect protection errors based on statically configured information on what the constituent parts of an OS-Application can do at runtime. See Section 7.7.

Unlike monitoring, protection facilities will trap the erroneous state at the point the error occurs, resulting in the shortest possible time between transition into an erroneous state and detection of the fault. The different kinds of protection errors are described in the glossary. If a protection error occurs before the Operating System module is started the behaviour is not defined. If a protection error happens during shutdown, e.g. in the application-specific shutdown hook, an endless loop between the shutdown service and the protection hook may occur.

In the case of a protection error, the Operating System module calls a user provided Protection Hook for the notification of protection errors at runtime. The Protection Hook runs in the context of the Operating System module and must therefore be trusted code.

The Operating System module itself needs only to detect an error and provide the ability to act. The Protection Hook can select one out of four options the Operating System module provides, which will be performed after returning from the Protection Hook, depending on the return value of the Protection Hook. The options are:

1. do nothing
2. forcibly terminate the faulty Task/Category 2 ISR
3. forcibly terminate all tasks and ISRs in the faulty OS-Application
 - a. without restart of the OS-Application
 - b. with restart of the OS-Application
4. shutdown the Operating System module.

Requirements [OS243](#) and [OS244](#) define the order of the default reaction if no faulty Task/Category 2 ISR or OS-Application can be found, e.g. in the system specific hook routines. Also OS-Applications are only mandatory in Scalability Classes 3 and 4, therefore in other Scalability Classes OS-Applications need not be defined.

Note that forcibly terminating interrupts is handled differently in “forcibly terminate the faulty ISR” and “forcibly terminate the OS-Application”. If a faulty ISR is forcibly terminated, the current invocation of the ISR is terminated. A subsequent invocation is allowed. If the OS-Application is forcibly terminated, then the interrupt source is also disabled, preventing subsequent interrupts.

The meaning of “do nothing” (protection hook returns `PRO_IGNORE`) is that the system ignores the error and continues operations as if no error happened at all. For example: a call to `ActivateTask()` causes an arrival rate violation and afterward the protection hook returns `PRO_IGNORE`. Then the actions to activate the task will be performed by the OS (maybe a rescheduling takes place) and the call returns with `E_OK` (if no other error was detected). The feature can be useful e.g. during development or to reach a higher availability (during development).

7.8.2 Requirements

[OS211] [The Operating System module shall execute the `ProtectionHook()` with the same permissions as the Operating System module.] ()

[OS107] [If no `ProtectionHook()` is configured and a protection error occurs, the Operating System module shall call `ShutdownOS()`.] (BSW11014)

[OS106] [If the `ProtectionHook()` returns `PRO_IGNORE` and was called with `E_OS_PROTECTION_ARRIVAL` the Operating System module shall return control to the user application.] (BSW11014)

[OS553] [If the `ProtectionHook()` returns `PRO_TERMINATETASKISR` the Operating System module shall forcibly terminate the faulty Task/Category 2 ISR.] ()

[OS554] [If the `ProtectionHook()` returns `PRO_TERMINATEAPPL` the Operating System module shall forcibly terminate the faulty OS-Application.] ()

[OS555] [If the `ProtectionHook()` returns `PRO_TERMINATEAPPL_RESTART` the Operating System module shall forcibly terminate the faulty OS-Application and afterwards restart the OS-Application.] ()

[OS556] [If the `ProtectionHook()` returns `PRO_SHUTDOWN` the Operating System module shall call the `ShutdownOS()`.] ()

[OS506] [If the `ProtectionHook()` is called with `E_OS_PROTECTION_ARRIVAL` the only valid return values are `PRO_IGNORE` or `PRO_SHUTDOWN`³. Returning other values will result in a call to `ShutdownOS()`.] ()

[OS475] [If the `ProtectionHook()` returns `PRO_IGNORE` and the `ProtectionHook()` was not called with `E_OS_PROTECTION_ARRIVAL` then the Operating System module shall call `ShutdownOS()`.] ()

[OS243] [If the `ProtectionHook()` returns `PRO_TERMINATETASKISR` and no Task or ISR can be associated with the error, the running OS-Application is forcibly terminated by the Operating System module. If even no OS-Application can be assigned, `ShutdownOS()` is called.] (BSW11014)

³ The reason for this case is that the Task which is supervised is not necessary active (and can not be e.g. terminated) and it can be that the caller of the activation is the real problem.

[OS244] [If the `ProtectionHook()` returns `PRO_TERMINATEAPPL` or `PRO_TERMINATEAPPL_RESTART` and no OS-Application can be assigned, `ShutdownOS()` is called.] (BSW11014)

[OS557] [If the `ProtectionHook()` returns `PRO_TERMINATEAPPL_RESTART` and no `OsRestartTask` was configured for the faulty OS-Application, `ShutdownOS()` is called.] ()

[OS108] [If the Operating System module forcibly terminates a task, it terminates the task, releases all allocated OSEK resources and calls `EnableAllInterrupts()/ResumeOSInterrupts() / ResumeAllInterrupts()` if the Task called `DisableAllInterrupts() / SuspendOSInterrupts() / SuspendAllInterrupts()` before without the corresponding `EnableAllInterrupts()/ResumeOSInterrupts() / ResumeAllInterrupts()` call.] (BSW11014)

[OS109] [If the Operating System module forcibly terminates an interrupt service routine, it clears the interrupt request, aborts the interrupt service routine (The interrupt source stays in the current state.) and releases all OSEK resources the interrupt service routine has allocated and calls `EnableAllInterrupts() / ResumeOSInterrupts() / ResumeAllInterrupts()` if the interrupt called `DisableAllInterrupts() / SuspendOSInterrupts() / SuspendAllInterrupts()` before without the corresponding `EnableAllInterrupts()/ResumeOSInterrupts() / ResumeAllInterrupts()` call.] (BSW11014)

[OS110] [If the Operating System module shall forcibly terminates an OS-Application, it:shall

- forcibly terminate all Tasks/ISRs of the OS-Application AND
- cancel all alarms of the OS-Application AND
- stop schedule tables of the OS-Application AND
- disable interrupt sources of Category 2 ISRs belonging to the OS-Application]

(BSW11014)

[OS111] [When the Operating System module restarts an OS-Application, it shall activate the configured `OsRestartTask`.] ()

7.9 Operating System for Multi-Core

This chapter specifies some extensions that allow to use an AUTOSAR system on Multi-Core micro-processors. It describes the main philosophy as well as additional

extensions to the existing OS functionality regarding Multi-Core. The following chapter contains a specification of a new mechanism within the OS called IOC (Inter OS-Application Communicator) that supports the communication between OS-Applications located on the same or on different cores

7.9.1 Background & Rationale

The existing AUTOSAR-OS is based on the OSEK/VDX Operating system which is widely used in the automotive industry. The AUTOSAR Multi-Core OS is derived from the existing AUTOSAR OS.

The Multi-Core OS in AUTOSAR is not a virtual ECU concept, instead it shall be understood as an OS that shares the same configuration and most of the code, but operates on different data structures for each core.

To reduce the memory footprint all cores should use the same code base.

Sometimes it can be beneficial to spend some more ROM/Flash, e.g. to use a local ROM, and "double" parts of the code to get faster ROM/Flash access.

7.9.1.1 Requirements

[OS567] [The generated part of the OS is derived from a single configuration that contains the relevant information for all cores. This implies, that IDs (e.g. TASKID, RESOURCEID, ...) are unique across cores. Every ID shall refer exactly to one entity independent from the core on which the entity is accessed. This applies also to objects that cannot be shared between cores.] (BSW4080008)

7.9.2 Scheduling

The priority of the TASKs drives the scheduling. Since multiple cores run truly parallel, several TASKs can execute at the same time.




Figure 2: Priorities are assigned to TASKs. The cores schedule independently from each other. The TASKs T₂, T₃ and T₅ are executed in true parallelism. TASKs with the same priority on the same core will be executed in order of activation; TASKs with the same priority on different

cores may not be executed in the order of activation, since the cores schedule independent from each other.

The OS can be entered on each core in parallel. This optimizes scalability towards multiple cores. The cores schedule independently. This implies that the schedule on one core does not consider the scheduling on the other cores⁴. A low priority TASK on one core may run in parallel with a high priority TASK on another core. TASKs and ISR's cannot dynamically change cores by means of the scheduling algorithm.

7.9.2.1 Requirements

[OS568] [Implementations shall be able to independently execute a TASK or an ISR on each started AUTOSAR OS core in parallel.] (BSW4080001)

[OS569] [The scheduling strategy as defined in AUTOSAR OS shall apply for each individual core in a Multi-Core system, for the TASKs and ISR assigned to the core.] (BSW4080001, BSW4080013)

7.9.3 Locatable entities (LE)

A locatable entity is an entity that has to be located entirely on one core. The assignment of LEs to cores is defined at configuration time (OsApplicationCoreAssignment).

In this release of the AUTOSAR standard OS-Applications shall be the LEs. Because every TASK has to run on some core, the usage of OS-Applications becomes obligatory in AUTOSAR R4.0 for Multi-Core systems. BSW modules are not allowed to ignore OS-Applications, even if they do not use any protection mechanisms. This is independent from the SC class.

As is stated in the AUTOSAR Specification of the Operating System, if OS-Applications are used, all Tasks, ISR etc. must belong to an OS-Application. This implies, that no AUTOSAR software exists outside of an OS-Application in Multi-Core systems.

On single-core systems OS-Applications are available only for SC3 and SC4 because the mechanism is used to support memory protection and implies the usage of extended mode. In Multi-core systems OS-Applications are always available independent of memory protection and on SC1 standard mode shall be possible.

7.9.3.1 Requirements

[OS570] [All TASKs that are assigned to the same OS-Application shall execute on the same core.] (BSW4080003, BSW4080005)

⁴ This also applies to TASKs with the same priority, bound to different cores. It also means that non-preemptive tasks cannot be preempted on the core they are running, but tasks on other cores can run in parallel.

[OS571] [All ISRs that are assigned to the same OS-Application shall execute on the same core.] (BSW4080003, BSW4080005)

[OS572] [ISR balancing (if supported by the HW) shall be switched off at boot time by the OS.] (BSW4080005, BSW4080006)

[OS764] [The OS module shall support OS-Applications in case of Multi-Core also for SC1 and SC2.] ()

[OS763] [In an SC1 system standard mode shall be possible.] ()

[OS573] [The binding of OS-Applications to cores shall be configured within the OS-Application container.] (BSW4080003, BSW4080005)

A new configuration item: `OsApplicationCoreAssignment {CORE}` within the OS-Application container shall be used to define the core to which the OS-Application is bound. The OS generator will map the configuration parameter "CORE" to a certain core, so that all OS-Applications with the same configuration parameter reside on the same core.

7.9.4 Multi-Core start-up concept

The way cores are started depends heavily on the hardware. Typically the hardware only starts one core, referred as the master core, while the other cores (slaves) remain in halt state until they are activated by the software.

In contrast to such a master-slave system other boot concepts with cores that start independently from each other are conceivable. However it is possible to emulate master-slave behavior on such systems by software.

The AUTOSAR Multi-Core OS specification requires a system with master-slave start-up behavior, either supported directly by the hardware or emulated in software. The master core is defined to be the core that requires no software activation, whereas a slave core requires activation by software.

In Multi-Core configurations, each slave core that is used by AUTOSAR must be activated before `StartOS` is entered on the core. Depending on the hardware, it may be possible to only activate a subset of the available cores from the master. The slave cores might activate additional cores before calling `StartOS`. All cores that belong to the AUTOSAR system have to be activated by the designated AUTOSAR API function. Additionally, the `StartOS` function has to be called on all these cores.

If a core is activated it executes some HW and compiler specific operations, before the "main" function is called. In case the same "main" function is executed on each core, the cores have to be differentiated by their specific core Id within the function.

Example:

```
void main ()
{
  StatusType rv;
  [...]
  switch (GetCoreID())
  {
    case OS_CORE_ID_MASTER:
      [...]
      StartCore(OS_CORE_ID_0, &rv);
      StartOS(OSDEFAULTAPPMODE);
      break;
    case OS_CORE_ID_0:
      [...]
      StartCore(OS_CORE_ID_1, &rv);
      StartOS(DONOTCARE);
      break;
    otherwise:
      StartOS(DONOTCARE);
  }
}
```

StartOS synchronizes all cores twice. The first synchronization point is located before the StartupHooks are executed, the second after the OS-Application specific StartupHooks have finished and before the scheduler is started. The exact point where the second synchronization occurs depends on the implementation, but it shall be before the scheduling is started. This release of the AUTOSAR specification does not support timeouts during the synchronization phase. Cores that are activated with StartCore but do not call StartOS may cause the system to hang. It is in the responsibility of the integrator to avoid such behavior.

As shown in Figure 3, the StartUpHook is called on every core right after the first synchronization. However, there is only one StartUpHook in the system. If, for example, core-individual functionality must be executed during StartUpHook the GetCoreID function can be used to discriminate the individual cores. After the global StartUpHook has finished each core performs the StartUpHooks of its OS-Applications. Since OS-Applications are bound to cores the OS-Application specific StartUpHooks are executed only on the core to which the corresponding OS-Application is bound.




Figure 3: This figure shows an example of an initialization process with 4 cores.

7.9.4.1 Requirements

[OS574] [The master core shall be able to activate cores.] (BSW4080006, BSW4080026, BSW4080027)

[OS575] [Any slave core shall be able to activate cores.] (BSW4080006, BSW4080026, BSW4080027)

[OS576] [It shall be allowed to use only a subset of the cores available on a μ C for the AUTOSAR system.] (BSW4080006, BSW4080026, BSW4080027)

[OS577] [The cores shall boot in master-slave mode. If this is not supported by the hardware, it shall be that the cores boot in parallel and emulate the behavior of a master-slave system.] (BSW4080006, BSW4080026, BSW4080027)

[OS578] [In case of an emulation a slave core (CoreS), which is controlled by the AUTOSAR OS (AUTOSAR core), shall not enter the main function before another core has activated the slave core by means of `StartCore(CoreS)`.] (BSW4080006)

[OS579] [All cores that belong to the AUTOSAR system shall be synchronized within the `StartOS` function before the scheduling is started and after the global `StartupHook` is called.] (BSW4080001, BSW4080006)

[OS580] [All cores that belong to the AUTOSAR system shall be synchronized within the `StartOS` before the global `StartupHook` is called.] (BSW4080006)

[OS581] [The global `StartupHook` shall be called on all cores immediately after the first synchronisation point.] (BSW4080006)

[OS582] [The OS-Application-specific StartupHooks shall be called after the global `StartupHook` but only on the cores to which the OS-Application is bound.]
(BSW4080006, BSW4080008)

7.9.5 Cores under control of the AUTOSAR OS

The AUTOSAR OS controls several cores as stated above. It need not control all cores of a μ C, however. The maximum number of controlled cores shall be configured within the “OsOS” section of the configuration.

The AUTOSAR OS API provides a `StartCore` function to start the cores under its control. The `StartCore` function takes a scalar value parameter of type `CoreIDType`, specifying the core that shall be started. `StartCore` can be called more than once on the master core and also on slave cores. Each core can only be started once, however. For example:

```
StartusType rv1, rv2;

StartCore(OS_CORE_ID_1, &rv1);
StartCore(OS_CORE_ID_2, &rv2);

if (rv1 != E_OK) || (rv2 != E_OK)
    EnterPanicMode();

StartOS(OSDEFAULTAPPMODE);
```

The `StartOS` function shall be called on all cores that have been activated by `StartCore`. It is not allowed to call `StartCore` from a core that has already called `StartOS`.

Cores that belong to the AUTOSAR system shall be started by the designated AUTOSAR OS API service `StartCore`.

7.9.5.1 Requirements

[OS583] [The number of cores that can be controlled by the AUTOSAR OS shall be configured offline.

A new configuration item (`OsNumberOfCores`) within the “**OsOS**” container is used to specify the maximum number of cores that are controlled by the AUTOSAR OS. If no value for (`OsNumberOfCores`) has been specified the number of cores shall be one.] (BSW4080001, BSW4080011)

7.9.6 Cores which are not controlled by the AUTOSAR OS

The function `StartNonAutosarCore` can be used both before and after `StartOS`. It is provided to activate cores that are controlled by another OS or no OS at all, AUTOSAR functions shall not be called on these cores, otherwise the behavior is unspecified.

7.9.6.1 Requirements

[OS584] [The AUTOSAR OS shall provide a function called `StartNonAutosarCore` that can be used to start cores, which are not controlled by the AUTOSAR OS.] (BSW4080006, BSW4080026, BSW4080027)

[OS585] [It shall be possible to activate cores that are not controlled by the AUTOSAR OS before and after calling `StartOS`.] (BSW4080006, BSW4080026, BSW4080027)

7.9.7 Multi-Core shutdown concept

AUTOSAR supports two shutdown concepts, the synchronized shutdown and the individual shutdown concept. While the synchronized shutdown is triggered by the new API function `ShutdownAllCores()`, the individual shutdown is invoked by the existing API function `ShutdownOS()`.

7.9.7.1 Synchronized shutdown concept

If a TASK with the proper rights calls “`ShutdownAllCores`”, a signal is sent to all other cores to induce the shutdown procedure. Once the shutdown procedure has started on a core, interrupts and TASKs are not further processed, and no scheduling will take place, therefore it makes no sense to activate any TASK, however no error will be generated. It is in the responsibility of the application developer/system integrator to make sure that any preparations for shutdown on application and basic software level are completed before calling “`ShutdownAllCores`”. (e.g. by means of the ECU state manager).

During the shutdown procedure every core executes its OS-Application specific `ShutdownHook` functions, followed by a synchronization point. After all cores have reached the synchronization point the global `ShutdownHook` function is executed by all cores in parallel.

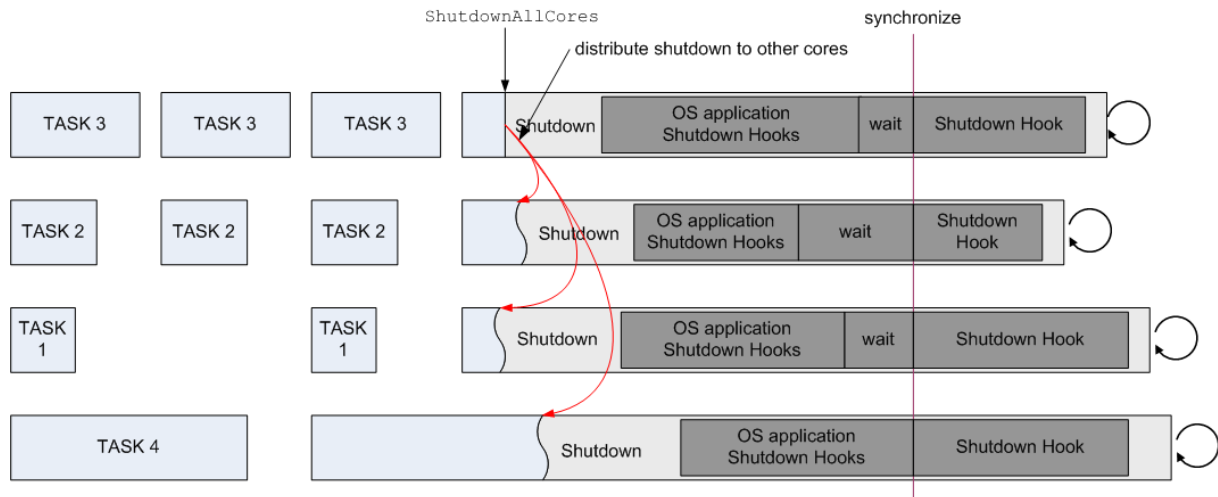


Figure 4: Example of a shutdown procedure.

[OS586] [During the shutdown, the OS-Application specific `ShutdownHook` shall be called on the core on which the corresponding OS-Application is bound.]
(BSW4080007)

[OS587] [Before calling the global `ShutdownHook`, all cores shall be synchronized.]
(BSW4080007)

[OS588] [The global `ShutdownHook` shall be called on all cores.] (BSW4080007)

7.9.7.2 Individual shutdown concept

If a TASK calls `ShutdownOS` the OS will be shut down on the core on which `ShutdownOS` has been called. Every core shall be able to invoke `ShutdownOS`. Similar to `StartOS` this function will shutdown the individual core. To shutdown the whole ECU `ShutdownOS` has to be called on every core. The function will not return. Individual shutdown is not supported in AUTOSAR R4.0 (AUTOSAR mode management will not use it).

7.9.7.3 Shutdown in case of fatal internal errors

In multicore systems it can happen that a fatal internal OS error is detected only on one core. In such cases a local shutdown of that core does not make sense.

[OS762] [In cases where the OS detects a fatal internal error all cores shall be shut down.] ()

7.9.8 OS service functionality (overview)

Within this chapter we describe which existing single core AUTOSAR OS functionality has been extended. The following table gives an overview of all standard OS API functions. The column “Multi-Core support” contains one of the following values:

- **Extended:** The function that has been extended substantially to support special Multi-Core functionality.
- **Adapted:** the function required some minor changes but basically remains unchanged.
- **Unchanged:** the behavior of the function has not changed.
- **New:** the function is a new AUTOSAR OS API-function.

Service	Multi-Core support	Annotation
ActivateTask	Extended	Cross core use shall be supported.
TerminateTask	Adapted	Check for unreleased spinlocks
ChainTask	Extended	Cross core use shall be supported.
Schedule	Adapted	Check for unreleased spinlocks
GetTaskID	Unchanged	Works only on the same core
GetTaskState	Extended	Cross core use shall be supported
DisableAllInterrupts	Unchanged	Works only on the same core
EnableAllInterrupts	Unchanged	Works only on the same core
SuspendAllInterrupts	Unchanged	Works only on the same core
ResumeAllInterrupts	Unchanged	Works only on the same core
SuspendOSInterrupts	Unchanged	Works only on the same core
ResumeOSInterrupts	Unchanged	Works only on the same core
GetResource	Adapted	Nestable with spinlocks
ReleaseResource	Adapted	Nestable with spinlocks
SetEvent	Extended	Cross core use shall be supported.
ClearEvent	Unchanged	
GetEvent	Unchanged	
WaitEvent	Adapted	Check for unreleased spinlocks
GetAlarmBase	Extended	Cross core use shall be supported
GetAlarm	Extended	Cross core use shall be supported
SetRelAlarm	Extended	Cross core use shall be supported
SetAbsAlarm	Extended	Cross core use shall be supported
CancelAlarm	Extended	Cross core use shall be supported
GetActiveApplicationMode	Unchanged	
StartOS	Extended	Support for MC systems
ShutdownOS	Extended	Support for MC systems
GetISRID	Unchanged	
GetApplicationID	Unchanged	
CallTrustedFunction	Adapted	Function must be bound to the same core
CheckISRMemoryAccess	Unchanged	
CheckTASKMemoryAccess	Unchanged	
CheckObjectAccess	Unchanged	
CheckObjectOwnership	Unchanged	
StartScheduleTableRel	Extended	Cross core use shall be supported.
StartScheduleTableAbs	Extended	Cross core use shall be supported.
StopScheduleTable	Extended	Cross core use shall be supported.
NextScheduleTable	Unchanged	
StartScheduleTableSynchron	Unchanged	
SyncScheduleTable	Unchanged	
GetScheduleTableStatus	Extended	Cross core use shall be supported.

SetScheduleTableAsync	Unchanged	
IncrementCounter	Adapted	Cross core is not allowed.
GetCounterValue	Extended	Cross core use shall be supported
GetElapsedCounterValue	Extended	Cross core use shall be supported.
TerminateApplication	Extended	Check for unreleased spinlocks. Cross core use shall be supported.
GetNumberOfActivatedCores	New	Number of cores activated during startup.
GetCoreID	New	ID of the current core
StartCore	New	Start additional core
StartNonAutosarCore	New	Start additional core
GetSpinlock	New	Occupy a spinlock
ReleaseSpinlock	New	Release a spinlock
TryToGetSpinlock	New	Try to occupy a spinlock
ShutdownAllCores	New	Synchronized shutdown.

Tab. 2: gives an overview of changes to the OS Service Calles

Service	Task	Cat1 ISR	Cat2 ISR	Error Hook	PreTask Hook	PostTask Hook	Startup Hook	Shutdown Hook	Alarm Callback	Protection Hook
GetNumberOfActivatedCores	✓		✓							
GetCoreID	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
StartCore										
StartNonAutosarCore										
GetSpinlock	✓		✓							
ReleaseSpinlock	✓		✓							
TryToGetSpinlock	✓		✓							
ShutdownAllCores	✓		✓	✓			✓			

Tab. 3: Allowed Calling Context for OS Service Calls

[OS589] [All functions that are not allowed to operate cross core shall return E_OS_CORE in extended status if called with parameters that require a cross core operation.] (BSW4080013)

7.9.9 GetTaskID

GetTaskID can be called both from TASK and ISR2 level. When called from an interrupt routine, on Single-Core systems, GetTaskID returns either the interrupted TASK or indicates that no TASK is running. On Multi-Core systems it

1. indicates that no TASK is running on the core or,
2. returns the ID of the interrupted TASK on the core.

7.9.10 Interrupt disabling

Note: All types of interrupts can only be disabled on the local core. This implies that the interrupt flags on other cores remain in their current state. Scheduling continues on the other cores. Running ISRs on other cores continue executing.

7.9.10.1 Requirements

[OS590] [The OS service "DisableAllInterrupts" shall only affect the core on which it is called.] (BSW4080013)

[OS591] [The OS service "EnableAllInterrupts" shall only affect the core on which it is called.] (BSW4080013)

[OS592] [The OS service "SuspendAllInterrupts" shall only affect the core on which it is called.] (BSW4080013)

[OS593] [The OS service "ResumeAllInterrupts" shall only affect the core on which it is called.] (BSW4080013)

[OS594] [The OS service "SuspendOSInterrupts" shall only affect the core on which it is called.] (BSW4080013)

[OS595] [The OS service "ResumeOSInterrupts" shall only affect the core on which it is called.] (BSW4080013)

7.9.11 TASK activation

TASK activation shall be extended to work across cores. This document will not specify any implementation details. This functions timing behavior can be slower when working across cores. If a TASK has to be activated on another core, a scheduling decision is necessary on that core. If the core has not been started an error is generated.

7.9.11.1 Requirements

[OS596] [It shall be possible to activate a TASK that is part of an OS-Application located on another core, as long as the assigned access rights allow it.] (BSW4080001, BSW4080015)

[OS598] [The call of `ActivateTask` across cores shall behave synchronously, i.e. a call returns after the task has been activated or an error has been detected. It shall

not be possible to continue execution on the calling core before `ActivateTask` is accomplished on the remote core.] (BSW4080015)

[OS599] [In case of an error when calling `ActivateTask` across cores, the error handler shall be called on the core on which `ActivateTask` was originally called.] (BSW4080015)

7.9.12 TASK Chaining

TASK chaining shall be extended to work across cores. This document will not specify any implementation details. This function's timing behavior can be slower when working across cores. If a TASK has to be activated on another core, a scheduling decision is necessary on that core. If the core has not been activated, an error is generated.

7.9.12.1 Requirements

[OS600] [It shall be possible to chain a TASK that is part of an OS-Application located on another core, as long as the assigned access rights allow it.] (BSW4080001, BSW4080015)

7.9.13 EVENT setting

`SetEvent` shall be extended to work across cores. This document will not specify any implementation details. This function's timing behavior can be slower when working across cores. If the core has not been activated, an error is generated.

7.9.13.1 Requirements

[OS602] [It shall be possible to set an EVENT that is part of an OS-Application located on another core, as long as the assigned access rights allow it.] (BSW4080016)

[OS604] [The call of `SetEvent` across cores shall behave synchronously, i.e. a call returns after the Event has been set or an error has been detected. It shall not be possible to continue execution on the calling core before `SetEvent` is accomplished on the remote core.] (BSW4080016)

[OS605] [In case of an error when calling `SetEvent` across cores, the error handler shall be called on the core on which `SetEvent` was originally called.] (BSW4080016)

7.9.14 Activating additional cores

The mechanism by which additional cores can be activated as described in section 7.9.5

7.9.15 Start of the OS

It is necessary to extend the functionality of `StartOS`. This is because `StartOS` is called once on each core. The user provides the so called application mode⁵ to the Operating System through the call parameter of `StartOS(AppMode)`. The application mode defines which of the configured (startup) objects (Tasks, Alarms, ScheduleTables) the OS automatically starts.

On a Multi-Core system all cores shall run in the same application mode. If `StartOS` is called with the Appmode `DONOTCARE`, the AppMode of the other cores is used. At least one core has to define an AppMode other than `DONOTCARE`.

If the application mode is the same on all cores, `StartOS` will proceed its task. More details can be found in chapter 7.9.4.

7.9.15.1 Requirements

[OS606] [The AUTOSAR specification does not support the activation of AUTOSAR cores after calling `StartOS` on that core. If `StartCore` is called after `StartOS` it shall return with `E_OS_ACCESS` in extended status.] (BSW4080001)

[OS607] [`StartOS` shall start the OS on the core on which it is called.] (BSW4080006, BSW4080013)

[OS608] [If more than one core calls `StartOS` with an AppMode other than "DONOTCARE", the AppModes shall be the same. `StartOS` shall check this at the first synchronisation point. In case of violation, `StartOS` shall not start the scheduling, shall not call any `StartupHooks`, and shall enter an endless loop on every core.] (BSW4080006)

[OS609] [If `StartOS` is called with the AppMode "DONOTCARE" the application mode of the other core(s) (differing from "DONOTCARE") shall be used.] (BSW4080006)

[OS610] [At least one core shall define an AppMode other than "DONOTCARE".] (BSW4080006)

⁵ This is the application mode of the Operating System and shall not be confused by other application modes defined in the AUTOSAR mode management.

[OS611] [If the IOC is configured, StartOS shall initialize the data structures of the IOC.] (BSW4080020)

7.9.16 TASK termination

The termination of TASKs requires an additional check: It is not allowed to terminate a TASK while a spinlock is occupied. If `TerminateTask / ChainTask` is called with an occupied spinlock an error is returned.

7.9.16.1 Requirements

If `TerminateTask` (or `ChainTask`) is called while the calling TASK holds a spinlock, the behavior is undefined in standard status.

[OS612] [In extended status `TerminateTask / ChainTask` shall return with an error (`E_OS_SPINLOCK`), which can be evaluated in the application.] (BSW4080021)

[OS613] [Spinlocks occupied by TASKs that are terminated in response to a protection hook shall be automatically released. This applies also to the case in which an OS-Application is terminated.] (BSW4080021)

7.9.17 Termination of OS-Applications

Similar to TASKs an OS-Application cannot be terminated while any of its TASKs occupy a spinlock. In such cases, the lock is automatically released. To avoid an avalanche of error handling, no calls to the `ErrorHook` are made.

It might be possible that `TerminateApplication(A)` is called in parallel from different cores. The implementation has to support such a call pattern by executing the first arriving call of `TerminateApplication(A)` and ignoring any subsequent calls until the termination is completed.

7.9.17.1 Requirements

[OS614] [`TerminateApplication` shall check if any of the TASKs in the OS-Application have occupied a spinlock. If so, the spinlocks shall be released.] (BSW4080021)

[OS615] [If `TerminateApplication(A)` is called in parallel from different cores, the OsApplication “A” is terminated by the first call, any subsequent calls will return with 'E_OK' in standard and extended status without doing anything, until the termination is completed.] (BSW4080021)

7.9.18 Shutdown of the OS

Every core shall be able to invoke shutdown by using the `ShutdownOS` function. By calling `ShutdownOS` only the calling core will enter the shutdown procedure.

If the user wants to shutdown all cores (more or less in parallel) `ShutdownAllCores` shall be used.

`ShutdownOS` and `ShutdownAllCores` will not return.

The OS service `ShutdownOS` is not used by the AUTOSAR mode management in AUTOSAR R4.0. The function is offered for users that run the OS on cores without RTE and without mode management.

7.9.18.1 Requirements

[OS616] [`ShutdownOS` shall be callable from each core running an AUTOSAR OS.] (BSW4080001, BSW4080007)

[OS617] [`ShutdownOS` shall shutdown the core on which it was called.] (BSW4080007)

[OS618] [The OS shall not start TASKs of an OS-Application once the shutdown procedure has been entered on a particular core.] (BSW4080013)

[OS619] [The AUTOSAR OS function `ShutdownOS` shall be callable in parallel on multiple cores.] (BSW4080013)

[OS620] [`ShutdownOS` shall release all spinlocks which are occupied by the calling core.] (BSW4080021)

[OS621] [`ShutdownAllCores` shall be callable from each core running an AUTOSAR OS.] (BSW4080007)

7.9.19 Waiting for EVENTS

The EVENT waiting mechanism must be adapted to the new Multi-Core spinlock functionality:

A TASK might be de-scheduled when calling `WaitEvent`, in which case it would not be able to release the spinlock. `WaitEvent` must therefore check if the calling TASK holds a spinlock. As with RESOURCES, spinlocks cannot be occupied by TASKs in wait state.

7.9.19.1 Requirements

[OS622] [The AUTOSAR Operating System `WaitEvent` API service shall check if it has been called while the calling TASK has occupied a spinlock. In extended status an error `E_OS_SPINLOCK` shall be returned and the TASK shall not enter the wait state.] (BSW4080021)

7.9.20 Calling trusted functions

Functions can be declared as trusted as part of an OS-Application. They can then only be executed through the `CallTrustedFunction` API function. Assuming that the access rights are configured accordingly, a TASK from OS-Application A can call a trusted function from OS-Application B.

On a Multi-Core system, these trusted function calls from one OS-Application to another are limited to the same core.

7.9.20.1 Requirements

[OS623] [The OS API function `CallTrustedFunction` shall return `E_OS_ACCESS` in extended status if the target trusted function is part of an OS-Application on another core.] (BSW4080013)

7.9.21 Invoking reschedule

The `Schedule` API service must be adapted to the new Multi-Core spinlock functionality in the same manner as `WaitEvent`.

A TASK shall not actively force a de-scheduling while it occupies spinlocks.

7.9.21.1 Requirements

[OS624] [The AUTOSAR Operating System `Schedule` API service shall check if it has been called while the calling TASK has occupied a spinlock. In extended status an error `E_OS_SPINLOCK` shall be returned and the scheduler shall not be called.] (BSW4080021)

7.9.22 RESOURCE occupation

The `GetResource` function allows mutual exclusion between TASKs on the same core. The OS generator shall check offline that the TASKs are not on different cores.(see 7.9.30) and the `GetResource` function will check this requirement online.

The priority ceiling protocol (used by `GetResource`) temporarily changes the priority of a `TASK`. Such an approach fails on Multi-Core systems as the priorities are local to each core. Therefore the ceiling protocol is not sufficient to protect a critical section against access from different cores.

7.9.23 The CoreID

Every HW assigns a unique physical Id to a core. The physical core Id is the only way to distinguish between cores. The physical core Ids of a μC are not necessarily consecutive and do not necessarily start with zero.

The SW requires a mechanism to identify a core, e.g. to use core specific variables. Because the physical core Id usually can not be used as a direct array index for core specific variables, a logical CoreID is necessary to map physical core Ids to array indexes. In the SW it is not necessary to know the physical core Id, the logical CoreID is sufficient.

The mapping of OSApplications and other SW objects to cores is specified in the configuration files. All such mappings shall be HW independent and therefore shall not be based on the physical core Id but on the logical CoreID.

The function `GetCoreID` internally maps the physical core Id to the logical CoreID. The mapping is implementation specific. `GetCoreID` can be either a C function or a macro.

7.9.23.1 Requirements

[OS625] [The AUTOSAR Operating System API function `GetCoreID` shall be callable before `StartOS`.] (BSW4080006)

[OS626] [An implementation shall offer a function `GetNumberOfActivatedCores` that returns the number of cores running the AUTOSAR OS.] (BSW4080001)

[OS627] [An implementation shall define a set of constants `OS_CORE_ID_<No>` of the type `CoreIDType` with `<No>` a value from 0 to "`OsNumberOfCores - 1`".] (BSW4080001)

[OS628] [An implementation shall offer a constant `OS_CORE_ID_MASTER` of the type `CoreIDType` that refers to the master core.] (BSW4080001)

7.9.24 COUNTERs, background & rationale

A COUNTER is represented by a COUNTER value, measured in “ticks”, and some COUNTER-specific constants.

Similarly to Single-Core situation, each operating system (on each core) offers at least one COUNTER that is derived from a timer. Therefore, it is possible to define several COUNTERs which belong to different OS-Applications and either resides on the same or different cores.

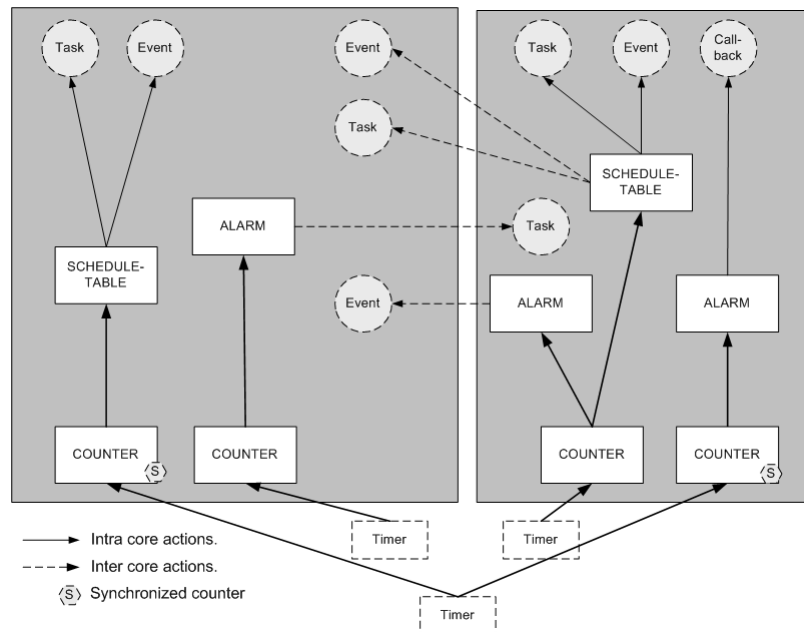


Figure 5: Examples of allowed configurations for COUNTERs, ALARMS, Schedule-tables and ISRs.

7.9.25 Multi-Core restrictions on COUNTERs

The AUTOSAR OS can only increment COUNTERSs on the core on which it resides. A COUNTER which is assigned to an OS-Application X cannot be incremented by an OS-Application Y if X and Y are assigned to different cores.

7.9.25.1 Requirements

[OS629] [A COUNTER belonging to an OS-Application shall be incremented by the core on which the OS-Application resides. The COUNTER shall not be incremented by other cores.] (BSW4080013)

[OS630] [It shall not be allowed to drive a schedule table from a COUNTER, which is assigned to a different core.] (BSW4080013)

[OS631] [It shall not be allowed to drive an ALARM from a COUNTER, which is assigned to a different core.] (BSW4080013)

There are two different reasons for these restrictions:

1. Race conditions can occur when cross-core modification of COUNTER is allowed (one core waits for a COUNTER to be modified by another core).
2. The core which is incrementing the COUNTER has to check if ALARMS which are based on the COUNTER have expired. Handling of expired ALARMS is more complex when different cores manipulate the same ALARMS, because mutual exclusion becomes necessary.

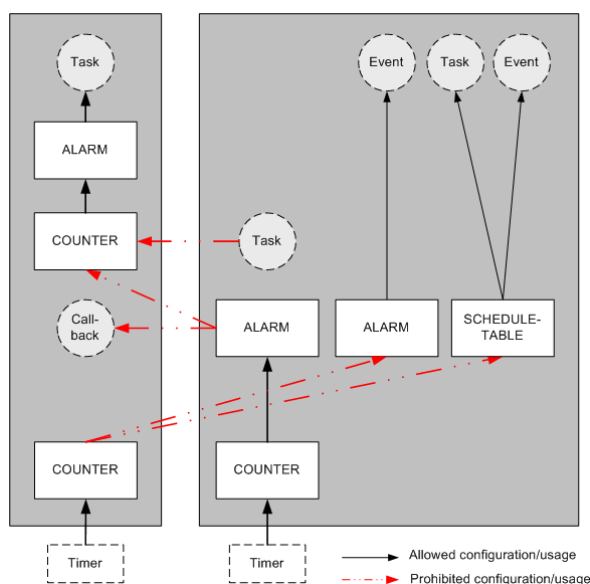


Figure 6: Example of disallowed configurations for COUNTERS, ALARMS, Schedule-tables and Call-backs.

7.9.26 Synchronization of COUNTERS

COUNTERS are used to drive ALARMS and schedule tables. To synchronize ALARMS and schedule tables that reside on different cores, the corresponding COUNTERS have to be synchronized.

For example, if the hardware supports this, it is possible that corresponding free running hardware counters on different cores use the same timer (same counter value maintained by the peripheral) and therefore provide the same timebase on different cores. Software COUNTERS can then get advanced by alarms attached to these core local corresponding hardware counters, e.g to drive synchronized schedule tables on different cores.

The quality of the synchronicity depends on the hardware architecture and on the system configuration. .

7.9.27 ALARMS

The ALARM mechanism of the AUTOSAR Operating System provides services to activate TASKs, set EVENTs, increment COUNTERs, or call an ALARM-call-back.

As stated above, ALARMS can only be bound to a COUNTER which resides on the same core. TASKs can be activated and EVENTs can be set with an ALARM action regardless of the core to which the TASK is bound. The access rights defined by OS-Applications have to be respected, however. Additionally it shall be allowed to manipulate ALARMS when they are bound to other cores. The API-services `SetRelAlarm`, `SetAbsAlarm`, and `CancelAlarm` can be used to manipulate parameters of ALARMS on other cores.

7.9.27.1 Requirements

[OS632] [If an ALARM expires, it shall be allowed to activate a TASK on a different core.] (BSW4080018)

[OS633] [If an ALARM expires, it shall be allowed to set an EVENT on a different core.] (BSW4080018)

[OS634] [The AUTOSAR Operating System shall process an ALARM on the core on which its corresponding OS-Application resides.] (BSW4080018)

[OS635] [ALARM callbacks shall be executed on the core to which the ALARM is bound. This is only applicable to SC1 systems, because otherwise Alarm Callback are not allowed (OS242).] (BSW4080013)

[OS636] [`SetRelAlarm` shall also work on an ALARM that is bound to another core.] (BSW4080013)

[OS637] [`SetAbsAlarm` shall also work on an ALARM that is bound to another core.] (BSW4080013)

[OS638] [`CancelAlarm` shall also work on an ALARM that is bound to another core.] (BSW4080013)

[OS639] [`GetAlarmBase` shall also work on an ALARM that is bound to another core.] (BSW4080013)

[OS640] [`GetAlarm` shall also work on an ALARM that is bound to another core.] (BSW4080013)

7.9.28 Schedule tables

Similarly to ALARMS, schedule tables can be used to activate TASKs and set EVENTs. As with ALARMS, a schedule table can only be bound to a COUNTER which resides on the same core.

To simplify system startup, it should be possible to start schedule tables on other cores. The system designer is responsible for the correct handling of schedule tables. For example, schedule tables can be controlled from one core.

7.9.28.1 Requirements

[OS641] [A schedule table shall be able to activate a TASK bound on a core other than the one upon which the schedule tables resides.] (BSW4080018)

[OS642] [A schedule table shall be able to set an EVENT on a core other than the one upon which the schedule tables resides] (BSW4080018)

[OS643] [The AUTOSAR Operating System shall process a schedule table on the core on which its corresponding OS-Application resides.] (BSW4080013)

[OS644] [The API call “StartScheduleTableAbs” shall be able to start schedule tables of OS-Applications residing on other cores.] (BSW4080018)

[OS645] [The API call “StartScheduleTableRel” shall be able to start schedule tables of OS-Applications residing on other cores.] (BSW4080013)

[OS646] [The API call “StopScheduleTable” shall be able to stop schedule tables of OS-Applications residing on other cores.] (BSW4080013)

[OS647] [The API service “GetScheduleTableStatus” shall be able to get the status of a schedule table that is part of an OS-Application residing on a different core.] (BSW4080013)

7.9.29 The spinlock mechanism

With the Multi-Core concept, a new mechanism is needed to support mutual exclusion for TASKs on different cores. This new mechanism shall not be used between TASKs on the same core because it makes no sense. In such cases the AUTOSAR Operating System returns an error.

A “SpinlockType”, which is similar to OSEK’s “ResourceType”, shall be used. Spinlocks are configured offline.

A spinlock is a busy waiting mechanism that polls a (lock) variable until it becomes available. Typically, this requires an atomic “test and set” functionality, the details of which are implementation specific.

Once a lock variable is occupied by a TASK/ISR2, other TASKs/ISR2s on other cores shall be unable to occupy the lock variable. The spinlock mechanism will not de-schedule these other TASKs while they poll the lock variable. However it might happen that a TASK/ISR with a higher priority becomes ready while the lock variable is being polled. In such cases the spinning TASK will be interfered. This is illustrated in Figure 7.

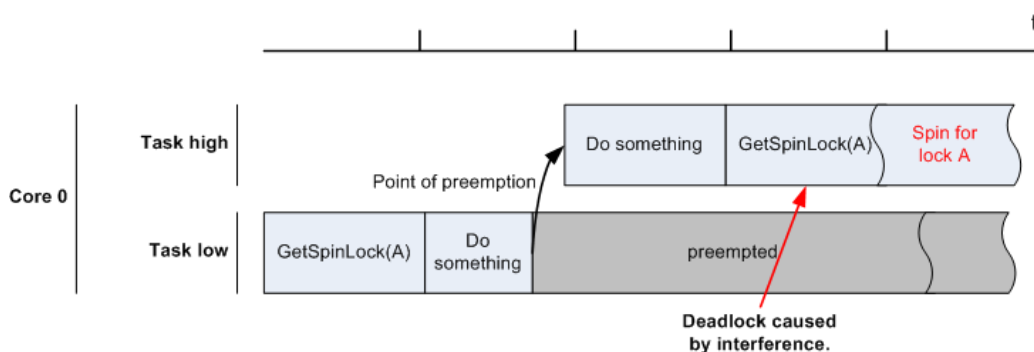


Figure 7: A deadlock situation caused by interference, the high priority TASK spins indefinitely because the low priority TASK has occupied the spinlock. In such cases the second GetSpinlock call will return with an error

A user can protect a TASK against such a situation by, for example, wrapping the spinlock with `SuspendAllInterrupts`, so that it cannot be interfered by other TASKs. A second deadlock situation can be created by nested spinlocks calls, as illustrated in Figure 8.

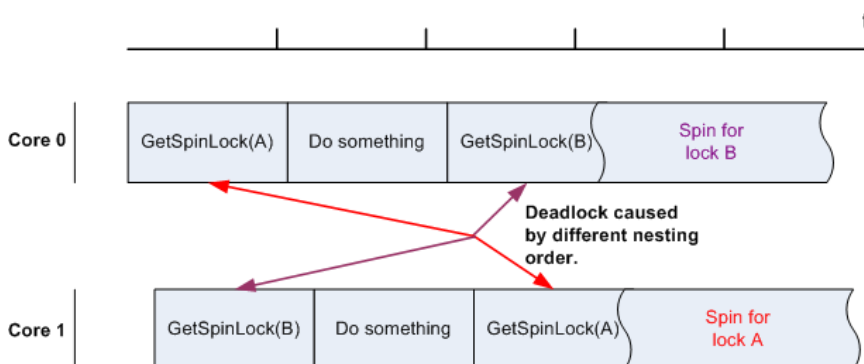


Figure 8: This figure shows a typical deadlock caused by two spinlocks taken in different order by TASKs on two different cores.

To avoid deadlocks it is not allowed to nest different spinlocks. Optionally if spinlocks shall be nested, a unique order has to be defined. Spinlocks can only be taken in this order whereas it is allowed to skip individual spinlocks. Cycles are not allowed within the defined order. This is illustrated in Figure 9.

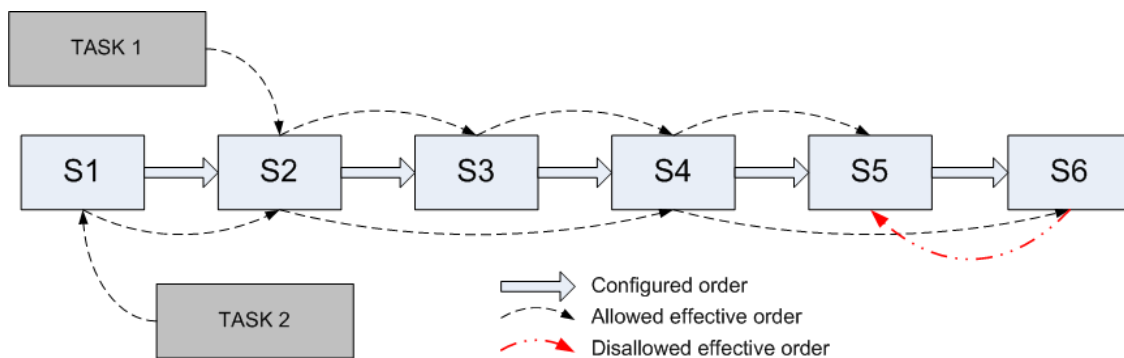


Figure 9: This figure shows an example in which two TASKS have access to a set of spinlocks S1 – S6. It is allowed to occupy the spinlocks in the predefined order and it is allowed to skip spinlocks. If multiple spinlocks are occupied at the same time, locking and unlocking has to occur in strict LIFO order.

The spinlock mechanism is not deadlock free by itself. The order in which spinlocks from Tasks/ISRs are requested has to be mentioned in the configuration description. If a task occupies a spinlock, scheduling shall be restricted.

7.9.29.1 Requirements

[OS648] [The AUTOSAR Operating System shall provide a spinlock mechanism that works across cores.] (BSW4080018, BSW4080021)

[OS649] [The AUTOSAR Operating System shall provide a `GetSpinlock` function which occupies a spinlock. If the spinlock is already occupied, `GetSpinlock` shall keep on trying to occupy the spinlock until it succeeds.] (BSW4080018, BSW4080021)

[OS650] [`GetSpinlock` shall be callable from TASK level.] (BSW4080018, BSW4080021)

[OS651] [`GetSpinlock` shall be callable from ISR2 level.] (BSW4080021)

The behavior of `GetSpinlock` is undefined if called from a category 1 ISR

[OS652] [The AUTOSAR Operating System shall provide a `TryToGetSpinlock` function which occupies a spinlock. If the spinlock is already occupied by a TASK, `TryToGetSpinlock` shall return.] (BSW4080018, BSW4080021)

[OS653] [`TryToGetSpinlock` shall be callable from TASK level.] (BSW4080018, BSW4080021)

[OS654] [`TryToGetSpinlock` shall be callable from ISR2 level.] (BSW4080018, BSW4080021)

[OS655] [The AUTOSAR Operating System shall provide a `ReleaseSpinlock` function which releases an occupied spinlock. If the spinlock is not occupied an error shall be returned.] (BSW4080018, BSW4080021)

[OS656] [`ReleaseSpinlock` shall be callable from TASK level.] (BSW4080018, BSW4080021)

[OS657] [`ReleaseSpinlock` shall be callable from ISR2 level.] (BSW4080018, BSW4080021)

[OS658] [The AUTOSAR Operating System shall generate an error if a TASK tries to occupy a spinlock that is assigned to a TASK/ISR2 on the same core (including itself).] (BSW4080018, BSW4080021)

[OS659] [The AUTOSAR Operating System shall generate an error if an ISR2 tries to occupy a spinlock that is assigned to a TASK/ISR2 on the same core.] (BSW4080018, BSW4080021)

[OS660] [A unique order in which multiple spinlocks can be occupied by a TASK/ISR2 should be configurable in the AUTOSAR Operating System. This might be realized by the configuration item (`OsSpinlockSuccessor{NEXT_SPINLOCK}`) where "NEXT_SPINLOCK" refers to the consecutive spinlock. (See chapter 10.2.5)] (BSW4080018, BSW4080021)

[OS661] [The AUTOSAR Operating System shall generate an error if a TASK/ISR2 that currently holds a spinlock tries to seize another spinlock that has not been configured as a direct or indirect successor of the latest acquired spinlock (by means of the `OsSpinlockSuccessor` configuration parameter) or if no successor is configured.] (BSW4080018, BSW4080021)

7.9.30 Offline checks

AUTOSAR RESOURCES cannot be shared between TASKs/ISR2s on different cores. The OS generator has to check if a user tries to assign a RESOURCE to TASKs on different cores and stop the generation process with an error.

COUNTERS cannot be accessed from OS-Applications on different cores. The OS generator has to reject configurations that violate this rule.

The linked list of spinlocks must be free of cycles to allow correct nesting of spinlocks in order to prevent deadlocks.

The OS generator tool must check that an OSApplication does not get assigned to a non existing core. Additional checks at configuration time, e.g. by an AUTOSAR description editor are recommended.

7.9.30.1 Requirements

[OS662] The OS generator tool shall return with an error if it detects a RESOURCE referred to by any TASKs or ISR's assigned to different cores.

] (BSW4080021)

[OS663] [The OS generator tool shall return with an error if an ALARM is assigned to a COUNTER on a different core.] (BSW4080013)

[OS664] [The OS generator tool shall return with an error if a COUNTER on a different core shall be incremented as an ALARM action.] (BSW4080013)

[OS665] [The OS generator tool shall return with an error if a schedule table is assigned to a COUNTER on a different core.] (BSW4080013)

[OS666] [The OS generator tool shall return with an error if the linked list of spinlocks is not free of cycles.] (BSW4080021)

[OS667] [The OS generator tool shall check the assignement of OsApplications (including the tasks assigned to the OsApplication) to cores and return an error in case any of these cores does not exist.] (BSW4080005)

7.9.31 Auto start Objects

Before scheduling starts the AUTOSAR Operating System⁶ activates all auto-start objects that are configured. This mechanism shall work similar on a Multi-Core system. Before scheduling starts, the Multi-Core OS shall activate all configured auto-start objects on the respective core. Due to the fact that OS-Applications are defined as the locatable entity no further configuration container is required. Auto-start objects are already configured as part of an OS-Application.

7.9.31.1 Requirements

[OS668] [The AUTOSAR Operating System shall automatically activate all auto-start TASKs configured for the current AppMode, with respect to the core, before the initial start of the scheduling.] (BSW4080006)

⁶ StartOS
106 of 230

[OS669] [The AUTOSAR Operating System shall automatically activate all auto-start ALARMS configured for the current AppMode, with respect to the core, before the initial start of the scheduling.] (BSW4080006)

[OS670] [The AUTOSAR Operating System shall automatically activate all auto-start schedule tables configured for the current AppMode, with respect to the core, before the initial start of the scheduling.] (BSW4080006)

7.10 Inter-OS-Application Communicator (IOC)

7.10.1 Background & Rationale

IOC stands for Inter OS-Application Communicator.

The "IOC" is responsible for the communication between OS-Applications and in particular for the communication crossing core or memory protection boundaries. Its internal functionality is closely connected to the Operating System.

[OS671] [The IOC implementation shall be part of the Operating System

The IOC is a third type of communication, in addition to

- Intra OS-Application communication: Always handled within the RTE
- Inter ECU communication: Already available via well defined interfaces to the communication stack (COM)] (BSW4080020)

Memory protection boundaries are a characteristic of OS-Applications and special communication mechanisms are needed to cross them. Multi-Core systems may also need additional measures to make communication between cores safe.

All AUTOSAR software, both BSW and software components, must belong to an OS-Application (s. 7.9.3), but not necessarily to the same one. It is expected that the BSW will be trusted code, but it shall be defined as one or more OS-Applications.

The IOC provides communication services between OS-Applications and in particular over core boundaries in Multi-Core systems. Because the cross-core communication is always an inter-OS-Application communication, the two mechanisms are combined. An inter OS-Application communication may not necessarily require a cross core communication, however.

Communication between OS-Applications is expected to be more frequent than inter ECU communication. This would be the case when existing; closely related Software

Components and their runnable entities are distributed to two or more cores to increase system performance. Meeting timing constraints is expected to become more difficult, when runnables which have been designed to run on a single core are distributed over several cores.

In systems with only one core, the IOC can be omitted completely, if just one OS-Application is available, or if no OS-Application uses memory protection mechanisms.

7.10.2 IOC - General purpose

The IOC provides communication services which can be accessed by clients which need to communicate across OS-Application boundaries on the same ECU. The RTE uses IOC services to communicate across such boundaries. All communication must be routed through the RTE on sender (or client) and on receiver (or server) side.

Direct access to IOC services by clients other than the RTE is currently not supported, but possible, if the client (e.g. a CDD) provides a hand written or generated IOC Configuration Description as specified and specific callback functions if necessary. Only sender/receiver communication is supported however by the IOC.

Software Components and/or BSW modules located in the same OS-Application (and hence on the same core) should not communicate by invoking IOC services. This would be less efficient than communication via RTE only. However, in case of IOC supported N:1 communication, if not all of the senders and the receiver are in the same OS-Application the IOC must be used.

To keep the RTE as hardware independent as possible, all inter OS-Application and inter core communication mechanisms and implementation variants are encapsulated in the IOC. The IOC internal functionality is dependent on hardware architecture properties, in particular on the memory architecture.

The IOC has to guarantee data consistency in inter OS-Application and inter core (Multi-Core systems) communication, this means in particular:

- In queued communication the sequential order of communication operations shall remain unchanged. In the N:1 communication case, the order of the messages from the different sources is a property of the implementation.
- The content of all data sent in one communication operation shall remain unchanged, i.e. each communication operation shall be treated as atomic operation.

7.10.3 IOC functionality

7.10.3.1 Communication

The IOC provides sender-receiver (signal passing) communication only. The RTE (or adapted BSW modules in a future release of this specification) translates Client-Server invocations and response transmissions into Sender-Receiver communication.

Only 1:1 and N:1 communication are supported by the IOC. In the case of 1:N communication, the RTE generates multiple calls to IOC services.

The IOC allows the transfer of one data item per atomic communication operation. A data item can either be a value for atomic basic data types or a reference for complex data structures. The data structure must be implemented as a single memory block, however. This way the data item can be transmitted in one piece. The IOC does not need to know the internal data structure. The basic memory address and length (which can be calculated from the type of the data item) is sufficient. The IOC does, e.g., not support a conversion of endianness between cores.

Transferring more than one data item in one operation is also supported for 1:1 communication only. In this case several types and memory addresses have to be used by the IOC function. The advantage compared to sequential IOC calls is that mechanisms to open memory protection boundaries and to notify the receiver have to be executed just once. Additionally, all data items are guaranteed to be consistent, because they are transferred in one atomic operation.

The IOC provides both, unqueued (Last-is-Best, data semantics) or queued (First-In-First-Out, event semantics) communication operations. If present, the IOC internal queue has a configurable length.

Each atomic communication operation gets specified individually by its own description block in a Configuration Description with regard to sender, receiver, data type(s), notification, and queuing.

7.10.3.2 Notification

The IOC optionally notifies the receiver as soon as the transferred data is available for access on the receiver side, by calling a configured callback function which gets provided by the user of the communication.

A possible implementation is to trigger an interrupt (Cat. 2) mechanism to invoke the callback function from the ISR on receiver side, or to use a microcontroller supplied trap. The callback function shall be efficient and compact, because it is called from within the ISR.

In certain cases, it might not be necessary to trigger an ISR to notify the receiver. The IOC generator can then select the appropriate IOC internal notification method based on the hardware architecture and other constraints. This might be more efficient than an ISR for communication between OsApplications on the same core.

The notification might be handled completely by the client of the IOC, e.g. when the RTE calls the IOC send function, and then notifies the receiver side RTE that new data are available from the IOC. In this case, the IOC is not affected at all by the details of the notification mechanism.

In case such alternative solutions prove to be more efficient, the IOC internal notification might get removed in future AUTOSAR releases.

7.10.4 IOC interface

The interface between RTE and IOC shall be similar to the interface between Software Components and the RTE, i.e. by generating specific interfaces for each communication operation instead of providing a generic API.

This supports optimization methods (like function inlining or replacing function calls by macros) much better than standardized interfaces. Most of the optimization can be performed offline at code generation time instead of consuming valuable real-time resources.

There is a unique set of IOC service APIs (at least to send and receive data) for each data communication specified in the IOC Configuration Description. Each service API gets generated and can be identified by a unique Id for each data communication. In case of N:1 communication, each sender must use its own API.

The same IOC service API and hence the same 1:1 communication can get used by more than one runnable inside the same SWC both on sender and on receiver side. However, the IOC functions are not reentrant, because otherwise e.g. spinlock errors could occur in case the IOC uses spinlocks in Multi-Core systems. The same IOC API must therefore only be called sequentially. This is no problem, if all runnable entities are scheduled within the same TASK, otherwise the caller is responsible to guarantee that the same IOC API is not called again before it returns from a different invocation.

Software Components may access the IOC only via RTE. Only the RTE decides which communication services to use to support the communication needs of Software Components.

Direct access to IOC services by BSW modules is not supported, but allowed for CDDs and other modules, if unavoidable. The clients have to provide a hand written or generated IOC Configuration Description as specified. In case of notification of the receiver, a specific callback function has to be specified and provided by the client. Only sender/receiver communication is supported however by the IOC.

7.10.5 IOC internal structure

This section gives some hints on possible IOC implementation options.

The IOC may enter the privileged mode to cross the protection boundaries between OS-Applications. The IOC therefore has to be part of the OS. Note that functionality that is placed in the kernel context might be non-interruptible by TASKs or ISR2. The functionality can be interrupted by Cat1 ISRs, however.

The IOC send service writes data into a buffer located in a memory area which is shared with the receiving communication partners (This is one possible implementation example using shared memory). Depending on the hardware architecture and other constraints, different implementation options might be available within the IOC. These options shall be transparent to the client (RTE), however.

The IOC ensures data consistency, i.e. there is a protection against concurrent access to the same data from all senders and the receiver for protection against inconsistent behavior and data corruption. The implementation can be hardware dependent.

In systems with shared memory, there can be a specific communication buffer for each data item in a memory section which is shared between the sending and receiving OS-Applications.

If an IOC communication with event semantics (queued) is configured the length of the queue shall be defined.

7.10.6 IOC configuration and generation

Data element specific interfaces between RTE and IOC require extensive code generation. Instead of generating the IOC together with the RTE, a sequential code generation process is used, to separate generic RTE code generation and hardware dependent IOC code generation as much as possible. The following steps shall be performed:

- Step 1: Specify all information about the allocation of Software Components to OS-Applications and cores in the ECU Configuration Description file.
- Step 2: Generate the RTE. The RTE generator creates data element specific IOC services calls and the corresponding IOC Configuration Description blocks (XML format) to specify the communication relations for each data element.
- Step 3: Generate the IOC code, according to the IOC Configuration Description (Step 2) while considering the hardware description files. Additionally, generate a header file (loc.h) for inclusion in RTE.c to provide definitions, function prototypes and macros.

Each atomic communication has to be specified in the IOC Configuration Description in a standardized XML format. There is one description block per communication operation specifying:

- Unique identifier
- Data type(s)
- Sender properties
- Receiver properties
- Name of callback function on receiver side in case of notification.
- Whether communication is queued or unqueued (last is best)
- In case of queued communication: Length of the queue

For details see Chapter 10.3

For each inter-OS-Application communication, the RTE generator creates one or more calls to an IOC function to send or receive data, and adds a corresponding description block to the IOC Configuration Description.

There are possibly multiple sources which contribute to the IOC configuration (e.g., RTE, CDD). The main input will come from the RTE generator. Other sources for the IOC Configuration Description (not supported in this specification revision) might be BSW module configuration tools or non-AUTOSAR components, which are allowed to use BSW services.

In ECUs with only one OS-Application, the IOC Configuration Description can be omitted.

7.10.7 IOC integration examples

This section describes two typical use cases that show how the IOC can support communication between OS-Applications. In both examples the OS-Applications are located on different cores of a Multi-Core system.

7.10.7.1 Example 1 - 1:1 sender/receiver communication without notification

One Software Component sends data items in "EVENT" semantics (queued) to another Software Component located on a different core. A runnable entity on the receiver side is invoked periodically (e.g. by an ALARM) and receives the data via RTE (see Figure 10).

Because the communication crosses core boundaries, the RTE invokes the IOC to transfer the data from core 0 to core 1.

On the sending side, the

```
Rte_Send_<port>_<item> (... , <data>)
```

call is mapped to an

```
IocSend_<Id> (<data>)
```


call.

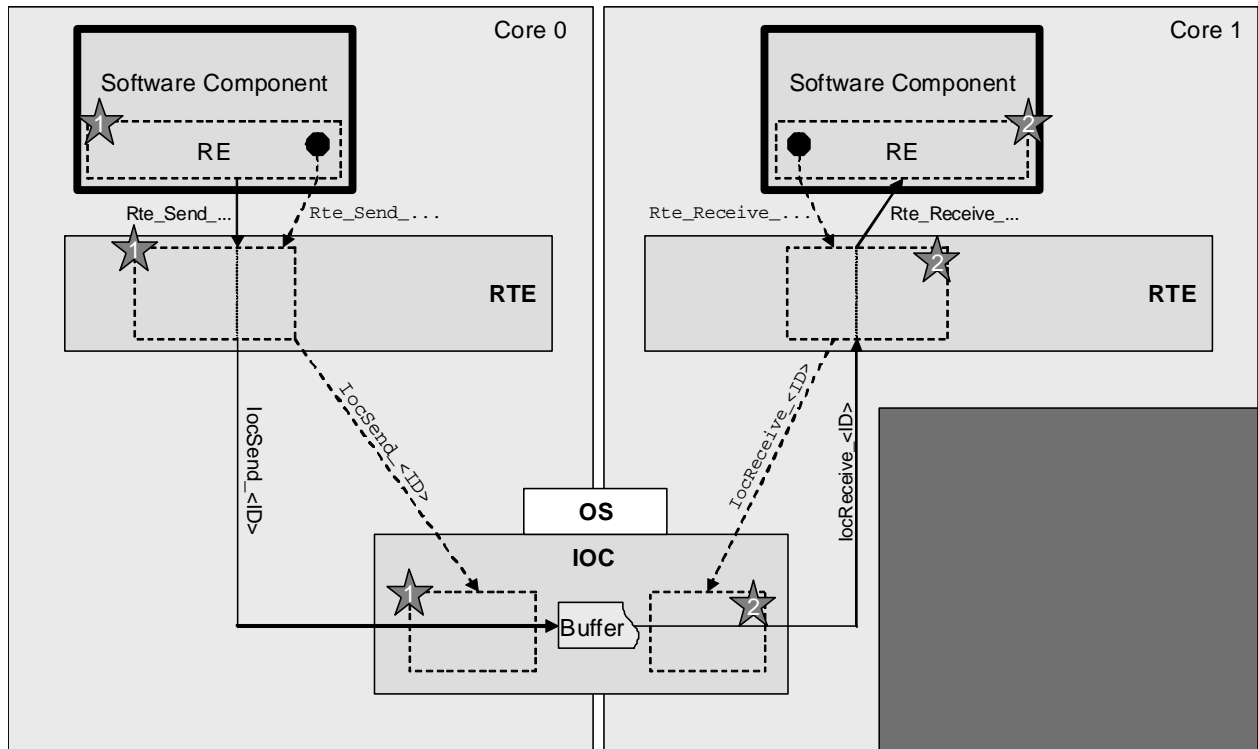


Figure 10: IOC without notification

In this example, the `IocSend` service writes the data into a buffer, located in a shared memory area which can get read by the receiver via the IOC.

On the receiving side, the receiving runnable gets invoked periodically. The

`Rte_Receive_<port>_<item> (... , <data>)`

call is mapped to an

`IocReceive_<Id> (<data>)`

call to read data from the IOC internal queue. An additional queue within the RTE is not necessary for 1:1 communication.

The IOC generator generates all the send and receive functions. The functions might be defined as macros for optimization purposes.

This kind of port to port communication without notification is suitable for:

- Sender/receiver communication
- Queued or unqueued communication
- 1:1 communication.

7.10.7.2 Example 2 - N:1 client/server communication with receiver notification by RTE

One Software Component invokes a service operation that is provided by another Software Component located on a different core. A runnable entity on the receiver side is activated to calculate the result (see Figure 11).

The RTE realizes the service on client side by mapping the client/server call to a sender/receiver communication. Because the communication crosses core boundaries, the RTE uses the IOC to transfer the data from Core 0 to Core 1.

On the sending side, the

```
Rte_Call_<port>_<op> (... , <data>)
```

call is mapped to a

```
IocSend_<Id> (<data>)
```

call to transmit the parameters over the IOC to the core hosting the server runnable.

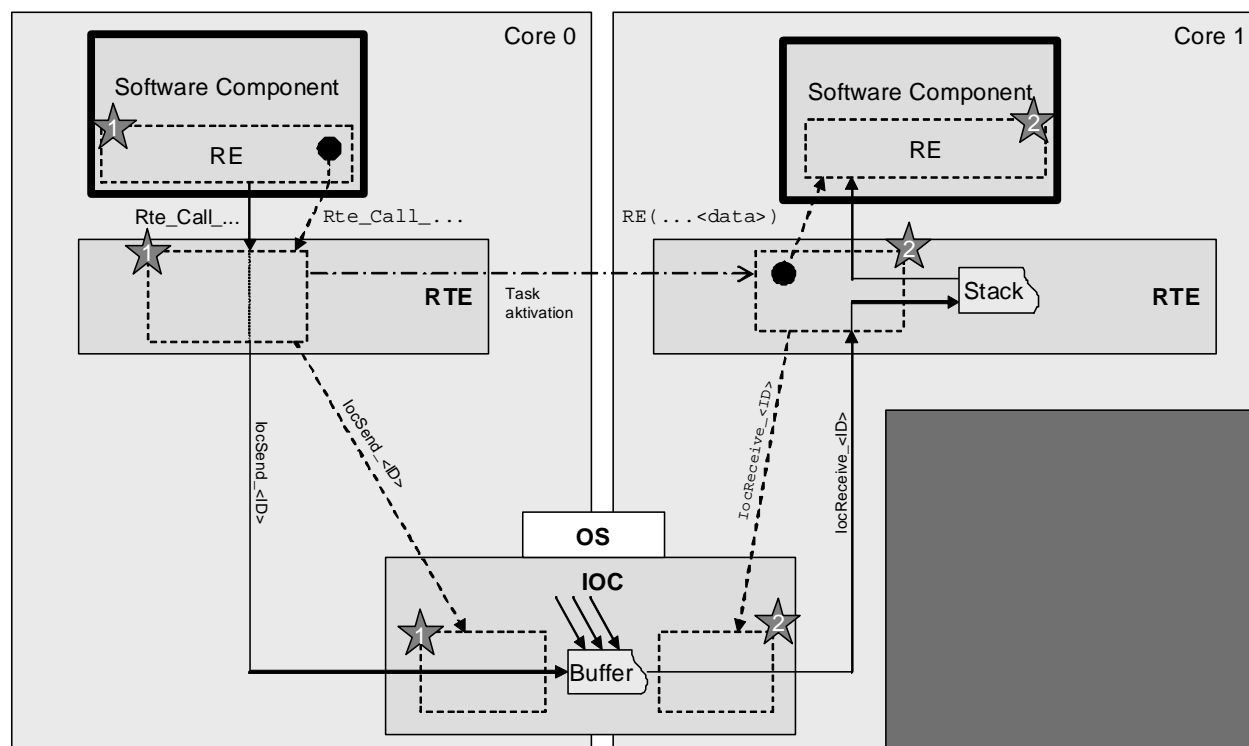


Figure 11: IOC with notification by RTE

After writing the data into the IOC internal queue buffer, the Rte_Call function uses an OS call to notify the receiver by activating the server TASK on the receiving core. This TASK is provided by the RTE. This TASK body is responsible for reading the data from the IOC buffer by calling IocReceive function and for forwarding the data to the server runnable. Depending on the return value of the IOC function, the IocReceive and server runnable calls might be repeated several times to empty the IOC internal queued buffer (if specified).

The result of the service on Core 1 is transferred back to the client on Core 0 in a similar way. The communication path of the result is not displayed in Figure 11.

This kind of port to port communication with notification by the RTE is suitable for:

- Sender/receiver communication with notification
- Client/server communication. In this case the RTE has to provide services to map the server call into 1:1 sender/receiver communication for the server call and another sender/receiver communication to return the result to the client
- Queued or unqueued communication
- 1:1 communication, if the receiver does not poll for data periodically (In this case, the solution in example 1 might have been more suitable)
- N:1 communication.

7.10.8 Future extensions

Some features are not supported by the first release of this specification, but might get added in a later release:

- In the future, the IOC will handle direct and efficient communication among BSW modules or between BSW modules and Software Components (via the RTE) located in different OS applications. Additional support of direct access from BSW modules to IOC services will be added.
- Other notification options (like activation of a specified TASK on receiver side) might be added later to the IOC.

7.11 System Scalability

7.11.1 Background & Rationale

In order to customize the operating system to the needs of the user and to take full advantage of the processor features the operating system can be scaled according to the following scalability classes

Feature	Described in Section	Scalability Class 1	Scalability Class 2	Scalability Class 3	Scalability Class 4	Hardware requirements
OSEK OS (all conformance classes)	7.1	✓	✓	✓	✓	
Counter Interface	8.4.16	✓	✓	✓	✓	
SWFRT Interface	8.4.17, 8.4.18	✓	✓	✓	✓	
Schedule Tables	7.3	✓	✓	✓	✓	
Stack Monitoring	7.5	✓	✓	✓	✓	
ProtectionHook	7.8		✓	✓	✓	
Timing Protection	7.7.2		✓		✓	Timer(s) with high priority interrupt
Global Time /Synchronization Support	7.4		✓		✓	Global time source
Memory Protection	7.7.1, 7.7.4			✓	✓	MPU
OS-Applications	7.6, 7.12			✓	✓	
Service Protection	7.7.3			✓	✓	
CallTrustedFunction	7.7.5			✓	✓	(Non-)privileged Modes

Tab. 4: Scalability classes

Feature	Scalability Class 1	Scalability Class 2	Scalability Class 3	Scalability Class 4
Minimum number of Schedule Tables supported	2	8	2	8
Minimum number of OS-Applications supported	0	0	2	2
Minimum number of software Counters supported	8	8	8	8

Tab. 5: Minimum requirements of scalability classes

7.11.2 Requirements

[OS240] [If an implementation of a lower scalability class supports features of higher classes then the interfaces for the features must comply with this Operating System software specification.] (BSW11012, BSW11016)

[OS241] [The Operating System module shall support the features according to the configured scalability class. (See Tab. 4)] (BSW11012, BSW11016)

[OS327] [The Operating System module shall always use extended status in Scalability Class 3 and 4.] ()

7.12 Hook Functions

7.12.1 Background & Rationale

Hook routines as defined in OSEK OS run at the level of the Operating System module and therefore can only belong to the trusted environment. Furthermore, these hook routines are global to the system (system-specific) and will probably be supplied by the ECU integrator.

In AUTOSAR however, each OS-Application may have the need to execute application specific code e.g. initialize some hardware in its own additional (application-specific) startup hook. These are called application specific hook routines. In general the application specific hooks have the same properties as the hook routines described in the OSEK OS specification. Differences are described below.

7.12.2 Requirements

[OS439] [The Operating System module shall provide the OSEK error macros (`OSError...()`) to all configured error hooks AND there shall be two (like in OIL) global configuration parameters to switch these macros on or off.] ()

StartupHook

[OS060] [If an application-specific startup hook is configured for an OS-Application `<App>`, the Operating System module shall call `StartupHook_<App>` on startup of the Operating System module.] ()

[OS226] [The Operating System module shall execute an application-specific startup hook with the access rights of the associated OS-Application.] ()

[OS236] [If both a system-specific and one (or more) application specific startup hook(s) are configured, the Operating System module shall call the system-specific startup hook before the application-specific startup hook(s).] ()

ShutdownHook

[OS112] [If an application-specific shutdown hook is configured for an OS-Application `<App>`, the Operating System module shall call `ShutdownHook_<App>` on shutdown of the OS.] ()

[OS225] [The Operating System module shall execute an application-specific shutdown hook with the access rights of the associated OS-Application.] ()

[OS237] [If both a system-specific and one (or more) application specific shutdown hook(s) are configured, the Operating System module shall call the system-specific shutdown hook after the application-specific shutdown hook(s).] ()

Error Hook

[OS246] [When an error occurs AND an application-specific error hook is configured for the faulty OS-Application <App>, the Operating System module shall call that application-specific error hook `ErrorHook_<App>` after the system specific error hook is called (if configured).] (BSW11013)

[OS085] [The Operating System module shall execute an application-specific error hook with the access rights of the associated OS-Application.] ()

[OS367] [Operating System module's services which do not return a `StatusType` shall not raise the error hook(s).] ()

7.13 Error classification

Instead of specifying two versions for production and development errors the AUTOSAR OS provides a finer granularity to adjust the error handling to specific needs, e.g. Scalability Classes, standard and extended status, switching on/off of hook routines.

<i>Type or error</i>	<i>Relevance</i>	<i>Related error code</i>	<i>Value</i>
Service can not be called.	Production	E_OS_SERVICEID	Assigned by implementation
An invalid address is given as a parameter to a service.	Production	E_OS_ILLEGAL_ADDRESS	Assigned by implementation
Tasks terminates without a <code>TerminateTask()</code> or <code>ChainTask()</code> call.	Production	E_OS_MISSINGEND	Assigned by implementation
A service of the OS is called inside an interrupt disable/enable pair.	Production	E_OS_DISABLEDINT	Assigned by implementation
A stack fault detected via stack monitoring by the OS	Production	E_OS_STACKFAULT	Assigned by implementation
A memory access violation occurred	Production	E_OS_PROTECTION_MEMORY	Assigned by implementation
A Task exceeds its execution time budget	Production	E_OS_PROTECTION_TIME	Assigned by implementation
A Category 2 ISR exceeds its execution time budget			

A Task/Category 2 ISR arrives before its timeframe has expired	Production	E_OS_PROTECTION_ARRIVAL	Assigned by implementation
A Task/Category 2 ISR blocks for too long	Production	E_OS_PROTECTION_LOCKED	Assigned by implementation
A trap occurred	Production	E_OS_PROTECTION_EXCEPTION	Assigned by implementation
Core is not available	Production	E_OS_CORE	Assigned by implementation
De-scheduling with occupied spinlock	Production	E_OS_SPINLOCK	Assigned by implementation
Deadlock situation due to interference	Production	E_OS_INTERFERENCE_DEADLOCK	Assigned by implementation
Potential deadlock due to wrong nesting	Production	E_OS_NESTING_DEADLOCK	Assigned by implementation

7.14 Debug support

In order to support debugging AUTOSAR implementations must publish information which can be used for debugging purpose.

[OS551] [Each variable that shall be accessible by AUTOSAR Debugging, shall be defined as global variable.] ()

[OS550] [All type definitions of variables which shall be debugged, shall be accessible by the header file Os.h.] ()

[OS549] [The declaration of variables in the header file shall be such, that it is possible to calculate the size of the variables by C-"sizeof".] ()

8 API specification

This chapter contains the APIs offered by the operating system. Note that not all services are available in all scalability classes, and that the behavior of some services is extended for specific scalability classes. For example, API to relatively start a schedule table has an additional check if the schedule table allows implicit synchronization. This check is only performed in SC2 and SC4 where synchronization of schedule tables is supported.

8.1 Constants

8.1.1 Error codes of type StatusType

The following constants are available in a multi-core environment.

Name:	AppModeType		
Type:	Enumeration		
Range:	DONOTCARE	--	
Description:	AppMode of the core shall be inherited from another core.		

Name:	TotalNumberOfCores		
Type:	scalar		
Range:	1..65535	--	--
Description:	The total number of cores		

Additional constants are in section 7.13 and [15].

8.2 Macros

```
OSMEMORY_IS_READABLE(<AccessType>)
OSMEMORY_IS_WRITEABLE(<AccessType>)
OSMEMORY_IS_EXECUTABLE(<AccessType>)
OSMEMORY_IS_STACKSPACE(<AccessType>)
```

These macros return a value not equal to zero if the memory is readable / writable / executable or stack space. The argument of the macros must be of type AccessType. Typically the return value of the service `Check[Task|ISR]MemoryAccess()` is used as argument for these macros.

8.3 Type definitions

8.3.1 ApplicationType (for OS-Applications)

Type:	Scalar
--------------	--------

Description:	This data type identifies the OS-Application.
Constants of this Type:	INVALID_OSAPPLICATION

8.3.2 ApplicationStateType

Type:	Scalar
Description:	This data type identifies the state of an OS-Application.
Constants of this Type:	APPLICATION_ACCESSIBLE APPLICATION_RESTARTING APPLICATION_TERMINATED

8.3.3 ApplicationStateRefType

Type:	Pointer
Description:	This data type points to location where a ApplicationStateType can be stored.

8.3.4 TrustedFunctionIndexType

Type:	Scalar
Description:	This data type identifies a trusted function.

8.3.5 TrustedFunctionParameterRefType

Type:	Pointer
Description:	This data type points to a structure which holds the arguments for a call to a trusted function.

8.3.6 AccessType

Type:	Integral
Description:	This type holds information how a specific memory region can be accessed.

8.3.7 ObjectAccessType

Type :	Scalar
Description:	This data type identifies if an OS-Application has access to an object.
Constants of this Type:	ACCESS NO_ACCESS

8.3.8 ObjectTypeType

Type :	Scalar
---------------	--------

Description :	This data type identifies an object.
Constants of this Type:	OBJECT_TASK OBJECT_ISR OBJECT_ALARM OBJECT_RESOURCE OBJECT_COUNTER OBJECT_SCHEDULETABLE

8.3.9 MemoryStartAddressType

Type:	Pointer
Description:	This data type is a pointer which is able to point to any location in the MCU address space.

8.3.10 MemorySizeType

Type:	Scalar
Description:	This data type holds the size (in bytes) of a memory region.

8.3.11 ISRType

Type:	Scalar
Description:	This data type identifies an interrupt service routine (ISR).
Constants of this Type:	INVALID_ISR

8.3.12 ScheduleTableType

Type:	Scalar
Description:	This data type identifies a schedule table.

8.3.13 ScheduleTableStatusType

Type:	Scalar
Description:	This type describes the status of a schedule. The status can be one of the following: <ul style="list-style-type: none"> ○ The schedule table is not started (SCHEDULETABLE_STOPPED) ○ The schedule table will be started after the end of currently running schedule table (schedule table was used in NextScheduleTable() service) (SCHEDULETABLE_NEXT) ○ The schedule table uses explicit synchronization, has been started and is waiting for the global time. (SCHEDULETABLE_WAITING) ○ The schedule table is running, but is currently not synchronous to a global time source (SCHEDULETABLE_RUNNING) ○ The schedule table is running and is synchronous to a global time source (SCHEDULETABLE_RUNNING_AND_SYNCHRONOUS)
Constants of this Type:	SCHEDULETABLE_STOPPED SCHEDULETABLE_NEXT SCHEDULETABLE_WAITING SCHEDULETABLE_RUNNING SCHEDULETABLE_RUNNING_AND_SYNCHRONOUS

8.3.14 ScheduleTableStatusRefType

Type:	Pointer
Description:	This data type points to a variable of the data type ScheduleTableStatusType.

8.3.15 CounterType

Type:	Scalar
Description:	This data type identifies a counter.

8.3.16 ProtectionReturnTypes

Type:	Scalar
Description:	This data type identifies a value which controls further actions of the OS on return from the protection hook.
Constants of this Type:	PRO_IGNORE PRO_TERMINATETASKISR PRO_TERMINATEAPPL PRO_TERMINATEAPPL_RESTART PRO_SHUTDOWN

8.3.17 RestartType

Type:	Scalar
Description:	This data type defines the use of a Restart Task after terminating an OS-Application.
Constants of this Type:	RESTART NO_RESTART

8.3.18 PhysicalTimeType

Type:	Scalar
Description:	This data type is used for values returned by the conversion macro (see OS393()) OS_TICKS2<Unit>_<Counter>().

8.3.19 CoreIDType

Name:	CoreIDType	
Type:	scalar	
Range:	OS_CORE_ID_MASTER	- refers to the master core, may be an alias for OS_CORE_ID_<x>
	OS_CORE_ID_0..OS_CORE_ID_65533	- refers to logical core 0, core 1 etc.
Description:	CoreIDType is a scalar that allows identifying a single core. The CoreIDType shall represent the logical CoreID	

8.3.20 SpinlockIdType

Name:	SpinlockIdType	
Type:	scalar	
Range:	INVALID_SPINLOCK	-- represents an invalid spinlock instance
	1..65535	-- 0x01, 0x02, ...: identifies a spinlock instance
Description:	SpinlockIdType identifies a spinlock instance and is used by the API functions: GetSpinlock, ReleaseSpinlock and TryToGetSpinlock.	

8.3.21 TryToGetSpinlockType

Name:	TryToGetSpinlockType	
Type:	Enumeration	
Range:	TRYTOGETSPINLOCK_SUCCESS	Spinlock successfully occupied
	TRYTOGETSPINLOCK_NOSUCCESS	Unable to occupy the spinlock
Description:	The TryToGetSpinlockType indicates if the spinlock has been occupied or not.	

8.4 Function definitions

The availability of the following services is defined in Tab. 4. The use of these services may be restricted depending on the context they are called from. See Tab. 1 for details.

8.4.1 GetApplicationID

[OS016] [

Service name:	GetApplicationID	
Syntax:	ApplicationType GetApplicationID(void)	
Service ID[hex]:	0x00	
Sync/Async:	Synchronous	
Reentrancy:	Reentrant	
Parameters (in):	None	
Parameters (inout):	None	
Parameters (out):	None	
Return value:	ApplicationType	<identifier of running OS-Application> or INVALID_OSAPPLICATION
Description:	This service determines the currently running OS-Application (a unique identifier has to be allocated to each application).	

] ()

[OS261] [GetApplicationID() shall return the application identifier to which the executing Task/ISR/hook belongs.] ()

[OS262] [If no OS-Application is running, GetApplicationID() shall return INVALID_OSAPPLICATION.] ()

[OS514] [Availability of GetApplicationID(): Available in Scalability Classes 3 and 4.] ()

8.4.2 GetISRID

[OS511] [

Service name:	GetISRID	
Syntax:	ISRType GetISRID(void)	
Service ID[hex]:	0x01	
Sync/Async:	Synchronous	
Reentrancy:	Reentrant	
Parameters (in):	None	
Parameters (inout):	None	
Parameters (out):	None	
Return value:	ISRType	<Identifier of running ISR> or INVALID_ISR
Description:	This service returns the identifier of the currently executing ISR.	

] ()

[OS263] [If called from category 2 ISR (or Hook routines called inside a category 2 ISR), GetISRID() shall return the identifier of the currently executing ISR.] ()

[OS264] [If its caller is not a category 2 ISR (or Hook routines called inside a category 2 ISR), GetISRID() shall return INVALID_ISR.] ()

[OS515] [Availability of `GetISRID()`: Available in all Scalability Classes.] ()

8.4.3 CallTrustedFunction

[OS097] [

Service name:	CallTrustedFunction	
Syntax:	<pre>StatusType CallTrustedFunction(TrustedFunctionIndexType FunctionIndex, TrustedFunctionParameterRefType FunctionParams)</pre>	
Service ID[hex]:	0x02	
Sync/Async:	Depends on called function. If called function is synchronous then service is synchronous. May cause rescheduling.	
Reentrancy:	Reentrant	
Parameters (in):	FunctionIndex	Index of the function to be called.
	FunctionParams	Pointer to the parameters for the function - specified by the FunctionIndex - to be called. If no parameters are provided, a NULL pointer has to be passed.
Parameters (inout):	None	
Parameters (out):	None	
Return value:	StatusType	E_OK: No Error E_OS_SERVICEID: No function defined for this index
Description:	A (trusted or non-trusted) OS-Application uses this service to call a trusted function	

] ()

[OS265] [If `<FunctionIndex>` is a defined function index, `CallTrustedFunction()` shall switch the processor into privileged mode AND shall call the function `<FunctionIndex>` out of a list of implementation specific trusted functions with disabled memory protection AND shall return `E_OK` after completion.]

()

[OS312] [Caveats of `CallTrustedFunction()`:

- The called trusted function must conform to the following C prototype: `void TRUSTED_<name_of_the_trusted_service>(TrustedFunctionIndexType,TrustedFunctionParameterRefType);` (The arguments are the same as the arguments of `CallTrustedFunction`).
- Normally, a user will not directly call this service, but it will be part of some standard interface, e.g. a standard I/O interface.
- It is the duty of the called trusted function to check rights of passed parameters, especially if parameters are interpreted as out parameters.
- It should be noted that the `CallTrustedFunction()` does not disable timing protection for the task which called the service. This may lead to timing faults (calls of the `ProtectionHook()`) even inside of a trusted OS-Application. It is therefore recommended to use `CallTrustedFunction()` only for stateless functions (e.g. functions which do not write or do not have internal states)] ()

[OS266] [When `CallTrustedFunction()` calls the function `<FunctionIndex>`, that function shall be executed with the same processor mode and memory protection boundaries as the OS-Application to which it belongs. It shall however retain the service protection limitations of the calling Task or ISR, and the notion of "current application" shall remain that of the calling Task or Category 2 ISR.] ()

Reaction to timing protection can be defined to terminate the OSApplication. If a task is inside `CallTrustedFunction()` and task rescheduling takes place within the same OSApplication, the newly running higher priority task may cause timing protection and terminate the OSApplication, thus indirectly aborting the trusted function. To avoid this, the scheduling of other Tasks which belong to the same OS-Application as the caller needs to be restricted, as well as the availability of interrupts of the same OS-Application.

[OS565] [When `CallTrustedFunction()` is called and the caller of `CallTrustedFunction()` is supervised with timing protection, the Operating System shall delay any timing protection errors until the return of `CallTrustedFunction().`] ()

[OS564] [If such a violation is detected inside a nested call sequence of `CallTrustedFunction()` of a task, the delay shall last until the return of the last `CallTrustedFunction().`] ()

[OS563] [The OperatingSystem shall not schedule any other Tasks which belong to the same OS-Application as the non-trusted caller of the service. Also interrupts of Category 2 which belong to the same OS-Application shall be disabled during the execution of the service.] ()

[OS364] [If `CallTrustedFunction()` calls the trusted function from a Category 2 ISR context, that function shall continue to run on the same interrupt priority and be allowed to call all system services defined for Category 2 ISR (see table in chapter 7.7.3.2).] ()

[OS365] [If `CallTrustedFunction()` calls the trusted function from a task context, that function shall continue to run on the same priority and be allowed to call all system services defined for tasks (see table in chapter 7.7.3.2).] ()

[OS292] [If the function index `<FunctionIndex>` in a call of `CallTrustedFunction()` is undefined, `CallTrustedFunction()` shall return `E_OS_SERVICEID.`] ()

[OS516] [Availability of `CallTrustedFunction()`: Available in Scalability Classes 3 and 4.] ()

8.4.4 CheckISRMemoryAccess

[OS512] [

Service name:	CheckISRMemoryAccess	
Syntax:	<pre>AccessType CheckISRMemoryAccess(ISRType ISRID, MemoryStartAddressType Address, MemorySizeType Size)</pre>	
Service ID[hex]:	0x03	
Sync/Async:	Synchronous	
Reentrancy:	Reentrant	
Parameters (in):	ISRID	ISR reference
	Address	Start of memory area
	Size	Size of memory area
Parameters (inout):	None	
Parameters (out):	None	
Return value:	AccessType	Value which contains the access rights to the memory area.
Description:	This service checks if a memory region is write/read/execute accessible and also returns information if the memory region is part of the stack space.	

] ()

[OS267] [If the ISR reference <ISRID> in a call of `CheckISRMemoryAccess()` is valid, `CheckISRMemoryAccess()` shall return the access rights of the ISR on the specified memory area.] ()

[OS313] [If an access right (e.g. “read”) is not valid for the whole memory area specified in a call of `CheckISRMemoryAccess()`, `CheckISRMemoryAccess()` shall yield no access regarding this right.] ()

[OS268] [If the ISR reference <ISRID> is not valid, `CheckISRMemoryAccess()` shall yield no access rights.] ()

[OS517] [Availability of `CheckISRMemoryAccess()`: Available in Scalability Classes 3 and 4.] ()

8.4.5 CheckTaskMemoryAccess

[OS513] [

Service name:	CheckTaskMemoryAccess	
Syntax:	<pre>AccessType CheckTaskMemoryAccess(TaskType TaskID, MemoryStartAddressType Address, MemorySizeType Size)</pre>	
Service ID[hex]:	0x04	
Sync/Async:	Synchronous	
Reentrancy:	Reentrant	

Parameters (in):	TaskID	Task reference
	Address	Start of memory area
	Size	Size of memory area
Parameters (inout):	None	
Parameters (out):	None	
Return value:	AccessType	Value which contains the access rights to the memory area.
Description:	This service checks if a memory region is write/read/execute accessible and also returns information if the memory region is part of the stack space.	

] ()

[OS269] [If the Task reference <TaskID> in a call of `CheckTaskMemoryAccess()` is valid, `CheckTaskMemoryAccess()` shall return the access rights of the task on the specified memory area.] ()

[OS314] [If an access right (e.g. “read”) is not valid for the whole memory area specified in a call of `CheckTaskMemoryAccess()`, `CheckTaskMemoryAccess()` shall yield no access regarding this right.] ()

[OS270] [If the Task reference <TaskID> in a call of `CheckTaskMemoryAccess()` is not valid, `CheckTaskMemoryAccess()` shall yield no access rights.] ()

[OS518] [Availability of `CheckTaskMemoryAccess()`: Available in Scalability Classes 3 and 4] ()

8.4.6 CheckObjectAccess

[OS256] [

Service name:	CheckObjectAccess	
Syntax:	<pre>ObjectAccessType CheckObjectAccess(ApplicationType ApplID, ObjectTypeType ObjectType, void ...)</pre>	
Service ID[hex]:	0x05	
Sync/Async:	Synchronous	
Reentrancy:	Reentrant	
Parameters (in):	ApplID	OS-Application identifier
	ObjectType	Type of the following parameter
	...	The object to be examined
Parameters (inout):	None	
Parameters (out):	None	
Return value:	ObjectAccessType	ACCESS if the ApplID has access to the object NO_ACCESS otherwise
Description:	This service determines if the OS-Applications, given by ApplID, is allowed to use the IDs of a Task, ISR, Resource, Counter, Alarm or Schedule Table in API calls.	

] ()

[OS271] [If the OS-Application <ApplID> in a call of `CheckObjectAccess()` has access to the queried object, `CheckObjectAccess()` shall return `ACCESS`.] ()

[OS272] [If the OS-Application <ApplID> in a call of `CheckObjectAccess()` has no access to the queried object, `CheckObjectAccess()` shall return `NO_ACCESS`.] ()

[OS423] [If in a call of `CheckObjectAccess()` the object to be examined is not a valid object OR <ApplID> is invalid OR <ObjectType> is invalid THEN `CheckObjectAccess()` shall return `NO_ACCESS`.] ()

[OS519] [Availability of `CheckObjectAccess()`: Available in Scalability Classes 3 and 4.] ()

8.4.7 CheckObjectOwnership

[OS017] [

Service name:	CheckObjectOwnership	
Syntax:	<pre>ApplicationType CheckObjectOwnership(ObjectTypeType ObjectType, void ...)</pre>	
Service ID[hex]:	0x06	
Sync/Async:	Synchronous	
Reentrancy:	Reentrant	
Parameters (in):	ObjectType	Type of the following parameter
	...	The object to be examined
Parameters (inout):	None	
Parameters (out):	None	
Return value:	ApplicationType	<OS-Application>: the OS-Application to which the object ObjectType belongs or <code>INVALID_OSAPPLICATION</code> if the object does not exists
Description:	This service determines to which OS-Application a given Task, ISR, Counter, Alarm or Schedule Table belongs	

] ()

[OS273] [If the object `ObjectType` specified in a call of `CheckObjectOwnership()` exists, `CheckObjectOwnership()` shall return the identifier of the OS-Application to which the object belongs.] ()

[OS274] [If in a call of `CheckObjectOwnership()` the specified object `ObjectType` is invalid OR the argument of the type (the "...") is invalid OR the object does not belong to any OS-Application, `CheckObjectOwnership()` shall return `INVALID_OSAPPLICATION`.] ()

[OS520] [Availability of CheckObjectOwnership():Available in Scalability Classes 3 and 4.] ()

8.4.8 StartScheduleTableRel

[OS347] [

Service name:	StartScheduleTableRel	
Syntax:	<pre>StatusType StartScheduleTableRel(ScheduleTableType ScheduleTableID, TickType Offset)</pre>	
Service ID[hex]:	0x07	
Sync/Async:	Synchronous	
Reentrancy:	Reentrant	
Parameters (in):	ScheduleTableID	Schedule table to be started
	Offset	Number of ticks on the counter before the the schedule table processing is started
Parameters (inout):	None	
Parameters (out):	None	
Return value:	StatusType	E_OK: No Error E_OS_ID (only in EXTENDED status): ScheduleTableID not valid. E_OS_VALUE (only in EXTENDED status): Offset is greater than (OsCounterMaxAllowedValue - InitialOffset) or is equal to 0. E_OS_STATE: Schedule table was already started.
Description:	This service starts the processing of a schedule table at "Offset" relative to the "Now" value on the underlying counter.	

] ()

[OS275] [If the schedule table <ScheduleTableID> in a call of StartScheduleTableRel() is not valid, StartScheduleTableRel() shall return E_OS_ID.] ()

[OS452] [If the schedule table <ScheduleTableID> in a call of StartScheduleTableRel() is implicitly synchronized (OsScheduleTblSyncStrategy = IMPLICIT), StartScheduleTableRel() shall return E_OS_ID.] ()

[OS332] [If <Offset> in a call of StartScheduleTableRel() is zero StartScheduleTableRel() shall return E_OS_VALUE.] ()

[OS276] [If the offset <Offset> is greater than OsCounterMaxAllowedValue of the underlying counter minus the Initial Offset, StartScheduleTableRel() shall return E_OS_VALUE.] ()

[OS277] [If the schedule table <ScheduleTableID> in a call of StartScheduleTableRel() is not in the state SCHEDULETABLE_STOPPED, StartScheduleTableRel() shall return E_OS_STATE.] ()

[OS278] [If the input parameters of StartScheduleTableRel() are valid and the state of schedule table <ScheduleTableID> is SCHEDULETABLE_STOPPED, then StartScheduleTableRel() shall start the processing of a schedule table <ScheduleTableID>. The Initial Expiry Point shall be processed after <Offset> + Initial Offset ticks have elapsed on the underlying counter. The state of <ScheduleTableID> is set to SCHEDULETABLE_RUNNING before the service returns to the caller.] ()

[OS521] [Availability of StartScheduleTableRel(): Available in all Scalability Classes.] ()

8.4.9 StartScheduleTableAbs

[OS358] [

Service name:	StartScheduleTableAbs	
Syntax:	<pre>StatusType StartScheduleTableAbs(ScheduleTableType ScheduleTableID, TickType Start)</pre>	
Service ID[hex]:	0x08	
Sync/Async:	Synchronous	
Reentrancy:	Reentrant	
Parameters (in):	ScheduleTableID	Schedule table to be started
	Start	Absolute counter tick value at which the schedule table is started
Parameters (inout):	None	
Parameters (out):	None	
Return value:	StatusType	E_OK: No Error E_OS_ID (only in EXTENDED status): ScheduleTableID not valid E_OS_VALUE (only in EXTENDED status): "Start" is greater than OsCounterMaxAllowedValue E_OS_STATE: Schedule table was already started
Description:	This service starts the processing of a schedule table at an absolute value "Start" on the underlying counter.	

] ()

[OS348] [If the schedule table <ScheduleTableID> in a call of StartScheduleTableAbs() is not valid, StartScheduleTableAbs() shall return E_OS_ID.] ()

[OS349] [If the <Start> in a call of StartScheduleTableAbs() is greater than the OsCounterMaxAllowedValue of the underlying counter, StartScheduleTableAbs() shall return E_OS_VALUE.] ()

[OS350] [If the schedule table <ScheduleTableID> in a call of StartScheduleTableAbs() is not in the state SCHEDULETABLE_STOPPED, StartScheduleTableAbs() shall return E_OS_STATE.] ()

[OS351] [If the input parameters of StartScheduleTableAbs() are valid and <ScheduleTableID> is in the state SCHEDULETABLE_STOPPED, StartScheduleTableAbs() shall start the processing of schedule table <ScheduleTableID> when the underlying counter next equals <Start> and shall set the state of <ScheduleTableID> to

- SCHEDULETABLE_RUNNING (for a non-synchronized / Explicitly synchronized schedule table) OR

- SCHEDULETABLE_RUNNING_AND_SYNCHRONOUS (for implicitly synchronized schedule table)

before returning to the user. (The Initial Expiry Point will be processed when the underlying counter next equals <Start>+Initial Offset).] ()

[OS522] [Availability of StartScheduleTableAbs(): Available in all Scalability Classes.] ()

8.4.10 StopScheduleTable

[OS006] [

Service name:	StopScheduleTable	
Syntax:	StatusType StopScheduleTable(ScheduleTableType ScheduleTableID)	
Service ID[hex]:	0x09	
Sync/Async:	Synchronous	
Reentrancy:	Reentrant	
Parameters (in):	ScheduleTableID	Schedule table to be stopped
Parameters (inout):	None	
Parameters (out):	None	
Return value:	StatusType	E_OK: No Error E_OS_ID (only in EXTENDED status): ScheduleTableID not valid. E_OS_NOFUNC: Schedule table was already stopped
Description:	This service cancels the processing of a schedule table immediately at any point while the schedule table is running.	

] ()

[OS279] [If the schedule table identifier <ScheduleTableID> in a call of StopScheduleTable() is not valid, StopScheduleTable() shall return E_OS_ID.] ()

[OS280] [If the schedule table with identifier <ScheduleTableID> is in state SCHEDULETABLE_STOPPED when calling StopScheduleTable(), StopScheduleTable() shall return E_OS_NOFUNC.] ()

[OS281] [If the input parameters of StopScheduleTable() are valid, StopScheduleTable() shall set the state of <ScheduleTableID> to SCHEDULETABLE_STOPPED and (stop the schedule table <ScheduleTableID> from processing any further expiry points and) shall return E_OK.] ()

[OS523] [Availability of StopScheduleTable(): Available in all Scalability Classes.] ()

8.4.11 NextScheduleTable

[OS191] [

Service name:	NextScheduleTable	
Syntax:	<pre>StatusType NextScheduleTable(ScheduleTableType ScheduleTableID_From, ScheduleTableType ScheduleTableID_To)</pre>	
Service ID[hex]:	0x0a	
Sync/Async:	Synchronous	
Reentrancy:	Reentrant	
Parameters (in):	ScheduleTableID_From	Currently processed schedule table
	ScheduleTableID_To	Schedule table that provides its series of expiry points
Parameters (inout):	None	
Parameters (out):	None	
Return value:	StatusType	E_OK: No error E_OS_ID (only in EXTENDED status): ScheduleTableID_From or ScheduleTableID_To not valid E_OS_NOFUNC: ScheduleTableID_From not started E_OS_STATE: ScheduleTableID_To is started or next
Description:	This service switches the processing from one schedule table to another schedule table.	

] (BSW099)

[OS282] [If the input parameter <ScheduleTableID_From> or <ScheduleTableID_To> in a call of NextScheduleTable() is not valid, NextScheduleTable() shall return E_OS_ID.] ()

[OS330] [If in a call of NextScheduleTable() schedule table <ScheduleTableID_To> is driven by different counter than schedule table <ScheduleTableID_From> then NextScheduleTable() shall return an error E_OS_ID.] ()

[OS283] [If the schedule table <ScheduleTableID_From> in a call of NextScheduleTable() is in state SCHEDULETABLE_STOPPED OR in state SCHEDULETABLE_NEXT, NextScheduleTable() shall leave the state of <ScheduleTable_From> and <ScheduleTable_To> unchanged and return E_OS_NOFUNC.] ()

[OS309] [If the schedule table <ScheduleTableID_To> in a call of NextScheduleTable() is not in state SCHEDULETABLE_STOPPED, NextScheduleTable() shall leave the state of <ScheduleTable_From> and <ScheduleTable_To> unchanged and return E_OS_STATE.] ()

[OS484] [If OsScheduleTblSyncStrategy of <ScheduleTableID_To> in a call of NextScheduleTable() is not equal to the OsScheduleTblSyncStrategy of <ScheduleTableID_From> then NextScheduleTable() shall return E_OS_ID.] ()

[OS284] [If the input parameters of NextScheduleTable() are valid then NextScheduleTable() shall start the processing of schedule table <ScheduleTableID_To> <ScheduleTableID_From>.FinalDelay ticks after the Final Expiry Point on <ScheduleTableID_From> is processed and shall return E_OK. NextScheduleTable() shall process the Initial Expiry Point on <ScheduleTableID_To> at <ScheduleTableID_From>.Final Delay + <ScheduleTable_To>.Initial Offset ticks after the Final Expiry Point on <ScheduleTableID_From> is processed.] ()

[OS324] [If the input parameters of NextScheduleTable() are valid AND the <ScheduleTableID_From> already has a “next” schedule table then NextScheduleTable() shall replace the previous “next” schedule table with <ScheduleTableID_To> and shall change the old “next” schedule table state to SCHEDULETABLE_STOPPED.] ()

[OS505] [If OsScheduleTblSyncStrategy of the schedule tables <ScheduleTableID_From> and <ScheduleTableID_To> in a call of NextScheduleTable() is EXPLICIT and the Operating System module already synchronizes <ScheduleTableID_From>, NextScheduleTable() shall continue synchronization after the start of processing <ScheduleTableID_To>.] ()

[OS453] [If the <ScheduleTableID_From> in a call of NextScheduleTable() is stopped, NextScheduleTable() shall not start the “next” schedule table and change its state to SCHEDULETABLE_STOPPED.] ()

[OS524] [Availability of NextScheduleTable(): Available in all Scalability Classes.] ()

8.4.12 StartScheduleTableSynchron

[OS201] [

Service name:	StartScheduleTableSynchron	
Syntax:	<pre>StatusType StartScheduleTableSynchron(ScheduleTableType ScheduleTableID)</pre>	
Service ID[hex]:	0x0b	
Sync/Async:	Synchronous	
Reentrancy:	Reentrant	
Parameters (in):	ScheduleTableID	Schedule table to be started
Parameters (inout):	None	
Parameters (out):	None	
Return value:	StatusType	E_OK: No Error E_OS_ID (only in EXTENDED status): ScheduleTableID not valid E_OS_STATE: Schedule table was already started
Description:	This service starts an explicitly synchronized schedule table synchronously.	

] (BSW11002)

[OS387] [If in a call of StartScheduleTableSynchron() the schedule table <ScheduleTableID> is not valid OR the schedule table <ScheduleTableID> is not explicitly synchronized (OsScheduleTblSyncStrategy != EXPLICIT) StartScheduleTableSynchron() shall return E_OS_ID.] ()

[OS388] [If the schedule table <ScheduleTableID> in a call of StartScheduleTableSynchron() is not in the state SCHEDULETABLE_STOPPED, StartScheduleTableSynchron() shall return E_OS_STATE.] ()

[OS389] [If <ScheduleTableID> in a call of StartScheduleTableSynchron() is valid, StartScheduleTableSynchron() shall set the state of <ScheduleTableID> to SCHEDULETABLE_WAITING and start the processing of schedule table <ScheduleTableID> after the synchronization count of the schedule table is set via SyncScheduleTable(). The Initial Expiry Point shall be processed when (Duration-SyncValue)+InitialOffset ticks have elapsed on the synchronization counter where:

- Duration is <ScheduleTableID>.OsScheduleTableDuration
- SyncValue is the <Value> parameter passed to the SyncScheduleTable()
- InitialOffset is the shortest expiry point offset in <ScheduleTableID>] ()

[OS525] [Availability of StartScheduleTableSynchron(): Available in Scalability Classes 2 and 4.] ()

8.4.13 SyncScheduleTable

[OS199] [

Service name:	SyncScheduleTable	
Syntax:	StatusType SyncScheduleTable(ScheduleTableType ScheduleTableID, TickType Value)	
Service ID[hex]:	0x0c	
Sync/Async:	Synchronous	
Reentrancy:	Reentrant	
Parameters (in):	ScheduleTableID	Schedule table to be synchronized
	Value	The current value of the synchronization counter
Parameters (inout):	None	
Parameters (out):	None	
Return value:	StatusType	E_OK: No errors E_OS_ID (only in EXTENDED status): The ScheduleTableID was not valid or schedule table can not be synchronized (OsScheduleTblSyncStrategy not set or OsScheduleTblSyncStrategy = IMPLICIT) E_OS_VALUE (only in EXETENDED status): The <Value> is out of range E_OS_STATE: The state of schedule table <ScheduleTableID> is equal to SCHEDULETABLE_STOPPED
Description:	This service provides the schedule table with a synchronization count and start synchronization.	

] (BSW11002)

[OS454] [If the <ScheduleTableID> in a call of SyncScheduleTable() is not valid OR schedule table can not be explicitly synchronized (OsScheduleTblSyncStrategy is not equal to EXPLICIT) SyncScheduleTable() shall return E_OS_ID.] ()

[OS455] [If the <Value> in a call of SyncScheduleTable() is greater than the OsScheduleTableDuration, SyncScheduleTable() shall return E_OS_VALUE.] ()

[OS456] [If the state of the schedule table <ScheduleTableID> in a call of SyncScheduleTable() is equal to SCHEDULETABLE_STOPPED or SCHEDULETABLE_NEXT SyncScheduleTable() shall return E_OS_STATE.] ()

[OS457] [If the parameters in a call of SyncScheduleTable() are valid, SyncScheduleTable() shall provide the Operating System module with the current synchronization count for the given schedule table. (It is used to synchronize the processing of the schedule table to the synchronization counter.)] ()

[OS526] [Availability of `SyncScheduleTable()`: Available in Scalability Classes 2 and 4.] ()

8.4.14 SetScheduleTableAsync

[OS422] [

Service name:	SetScheduletableAsync	
Syntax:	<pre>StatusType SetScheduletableAsync(ScheduleTableType ScheduleTableID)</pre>	
Service ID[hex]:	0x0d	
Sync/Async:	Synchronous	
Reentrancy:	Reentrant	
Parameters (in):	ScheduleTableID	Schedule table for which status is requested
Parameters (inout):	None	
Parameters (out):	None	
Return value:	StatusType	E_OK: No Error E_OS_ID (only in EXTENDED status): Invalid ScheduleTableID
Description:	This service stops synchronization of a schedule table.	

] ()

[OS362] [If `SetScheduleTableAsync()` is called for a running schedule table, the Operating System module shall stop further synchronization until a `SyncScheduleTable()` call is made.] ()

[OS323] [If `SetScheduleTableAsync()` is called for a running schedule table the Operating System module shall continue to process expiry points on the schedule table.] ()

[OS458] [If `OsScheduleTblSyncStrategy` of `<ScheduleTableID>` in a call of `SetScheduleTableAsync()` is not equal to `EXPLICIT` OR if `<ScheduleTableID>` is invalid then `SetScheduleTableAsync()` shall return `E_OS_ID`.] ()

[OS483] [If the current state of the `<ScheduleTableID>` in a call of `SetScheduleTableAsync()` equals to `SCHEDULETABLE_STOPPED`, `SCHEDULETABLE_NEXT` or `SCHEDULETABLE_WAITING` then `SetScheduleTableAsync()` shall return `E_OS_STATE`.] ()

[OS300] [If the current state of `<ScheduleTableID>` in a call of `SetScheduleTableAsync()` equals `SCHEDULETABLE_RUNNING_AND_SYNCHRONOUS` (or `SCHEDULETABLE_RUNNING`) then `SetScheduleTableAsync()` shall set (or keep in case of `SCHEDULETABLE_RUNNING`) the status of `<ScheduleTableID>` to `SCHEDULETABLE_RUNNING`.] ()

[OS527] [Availability of SetScheduleTableAsync(): Available in Scalability Classes 2 and 4.] ()

8.4.15 GetScheduleTableStatus

[OS227] [

Service name:	GetScheduleTableStatus	
Syntax:	<pre>StatusType GetScheduleTableStatus(ScheduleTableType ScheduleTableID, ScheduleTableStatusRefType ScheduleStatus)</pre>	
Service ID[hex]:	0x0e	
Sync/Async:	Synchronous	
Reentrancy:	Reentrant	
Parameters (in):	ScheduleTableID	Schedule table for which status is requested
Parameters (inout):	None	
Parameters (out):	ScheduleStatus	Reference to ScheduleTableStatusType
Return value:	StatusType	E_OK: No Error E_OS_ID (only in EXTENDED status): Invalid ScheduleTableID
Description:	This service queries the state of a schedule table (also with respect to synchronization).	

] (BSW11002)

[OS289] [If the schedule table <ScheduleTableID> in a call of GetScheduleTableStatus() is NOT started, GetScheduleTableStatus() shall pass back SCHEDULETABLE_STOPPED via the reference parameter <ScheduleStatus> AND shall return E_OK.] ()

[OS353] [If the schedule table <ScheduleTableID> in a call of GetScheduleTableStatus() was used in a NextScheduleTable() call AND waits for the end of the current schedule table, GetScheduleTableStatus() shall return SCHEDULETABLE_NEXT via the reference parameter <ScheduleStatus> AND shall return E_OK.] ()

[OS354] [If the schedule table <ScheduleTableID> in a call of GetScheduleTableStatus() is configured with explicit synchronization AND <ScheduleTableID> was started with StartScheduleTableSynchron() AND no synchronization count was provided to the Operating System, GetScheduleTableStatus() shall return SCHEDULETABLE_WAITING via the reference parameter <ScheduleStatus> AND shall return E_OK.] ()

[OS290] [If the schedule table <ScheduleTableID> in a call of GetScheduleTableStatus() is started AND synchronous, GetScheduleTableStatus() shall pass back SCHEDULETABLE_RUNNING_AND_SYNCHRONOUS via the reference parameter <ScheduleStatus> AND shall return E_OK.] ()

[OS291] [If the schedule table <ScheduleTableID> in a call of GetScheduleTableStatus() is started AND NOT synchronous (deviation is not within the precision interval OR the schedule table has been set asynchronous), GetScheduleTableStatus() shall pass back SCHEDULETABLE_RUNNING via the reference parameter ScheduleStatus AND shall return E_OK.] ()

[OS293] [If the identifier <ScheduleTableID> in a call of GetScheduleTableStatus() is NOT valid, GetScheduleTableStatus() shall return E_OS_ID.] ()

[OS528] [Availability of GetScheduleTableStatus(): Available in all Scalability Classes.] ()

8.4.16 IncrementCounter

[OS399] [

Service name:	IncrementCounter	
Syntax:	StatusType IncrementCounter(CounterType CounterID)	
Service ID[hex]:	0x0f	
Sync/Async:	Synchronous, may cause rescheduling	
Reentrancy:	Reentrant	
Parameters (in):	CounterID	The Counter to be incremented
Parameters (inout):	None	
Parameters (out):	None	
Return value:	StatusType	E_OK: No errors E_OS_ID (only in EXTENDED status): The CounterID was not valid or counter is implemented in hardware and can not be incremented by software
Description:	This service increments a software counter.	

] ()

[OS285] [If the input parameter <CounterID> in a call of IncrementCounter() is not valid OR the counter is a hardware counter, IncrementCounter() shall return E_OS_ID.] ()

[OS286] [If the input parameter of IncrementCounter() is valid, IncrementCounter() shall increment the counter <CounterID> by one (if any alarm connected to this counter expires, the given action, e.g. task activation, is done) and shall return E_OK.] (BSW11020)

[OS321] [If in a call of IncrementCounter() an error happens during the execution of an alarm action, e.g. E_OS_LIMIT caused by a task activation,

IncrementCounter() shall call the error hook(s), but the IncrementCounter() service itself shall return E_OK.] ()

[OS529] [Caveats of IncrementCounter(): If called from a task, rescheduling may take place.] ()

[OS530] [Availability of IncrementCounter(): Available in all Scalability Classes.] ()

8.4.17 GetCounterValue

[OS383] [

Service name:	GetCounterValue	
Syntax:	StatusType GetCounterValue(CounterType CounterID, TickRefType Value)	
Service ID[hex]:	0x10	
Sync/Async:	Synchronous	
Reentrancy:	Reentrant	
Parameters (in):	CounterID	The Counter which tick value should be read
Parameters (inout):	None	
Parameters (out):	Value	Contains the current tick value of the counter
Return value:	StatusType	E_OK: No errors E_OS_ID (only in EXTENDED status): The <CounterID> was not valid
Description:	This service reads the current count value of a counter (returning either the hardware timer ticks if counter is driven by hardware or the software ticks when user drives counter).	

] (SWFRT00025)

[OS376] [If the input parameter <CounterID> in a call of GetCounterValue() is not valid, GetCounterValue() shall return E_OS_ID.] ()

[OS377] [If the input parameter <CounterID> in a call of GetCounterValue() is valid, GetCounterValue() shall return the current tick value of the counter via <Value> and return E_OK.] (SWFRT00033)

[OS531] [Caveats of GetCounterValue(): Note that for counters of OsCounterType = HARDWARE the real timer value (the – possibly adjusted – hardware value, see [OS384](#)) is returned, whereas for counters of OsCounterType = SOFTWARE the current “software” tick value is returned.] ()

[OS532] [Availability of GetCounterValue(): Available in all Scalability Classes.] ()

8.4.18 GetElapsedValue

[OS392] [
Service name:	GetElapsedValue	
Syntax:	<pre>StatusType GetElapsedValue(CounterType CounterID, TickRefType Value, TickRefType ElapsedValue)</pre>	
Service ID[hex]:	0x11	
Sync/Async:	Synchronous	
Reentrancy:	Reentrant	
Parameters (in):	CounterID	The Counter to be read
Parameters (inout):	Value	in: the previously read tick value of the counter out: the current tick value of the counter
Parameters (out):	ElapsedValue	The difference to the previous read value
Return value:	StatusType	E_OK: No errors E_OS_ID (only in EXTENDED status): The CounterID was not valid E_OS_VALUE (only in EXTENDED status): The given Value was not valid
Description:	This service gets the number of ticks between the current tick value and a previously read tick value.	

] (SWFRT00025)

[OS381] [If the input parameter <CounterID> in a call of GetElapsedValue() is not valid GetElapsedValue() shall return E_OS_ID.] ()

[OS391] [If the <Value> in a call of GetElapsedValue() is larger than the max allowed value of the <CounterID>, GetElapsedValue() shall return E_OS_VALUE.] ()

[OS382] [If the input parameters in a call of GetElapsedValue() are valid, GetElapsedValue() shall return the number of elapsed ticks since the given <Value> value via <ElapsedValue> and shall return E_OK.] (SWFRT00034)

[OS460] [GetElapsedValue() shall return the current tick value of the counter in the <Value> parameter.] ()

[OS533] [Caveats of GetElapsedValue():If the timer already passed the <Value> value a second (or multiple) time, the result returned is wrong. The reason is that the service can not detect such a relative overflow.] ()

[OS534] [Availability of GetElapsedValue(): Available in all Scalability Classes.] ()

8.4.19 TerminateApplication

[OS258] [

Service name:	TerminateApplication	
Syntax:	StatusType TerminateApplication(ApplicationType Application, RestartType RestartOption)	
Service ID[hex]:	0x12	
Sync/Async:	Synchronous	
Reentrancy:	Reentrant	
Parameters (in):	Application	The identifier of the OS-Application to be terminated. If the caller belongs to <Application> the call results in a self termination.
	RestartOption	Either RESTART for doing a restart of the OS-Application or NO_RESTART if OS-Application shall not be restarted.
Parameters (inout):	None	
Parameters (out):	None	
Return value:	StatusType	E_OK: No errors E_OS_ID: <Application> was not valid (only in EXTENDED status) E_OS_VALUE: <RestartOption> was neither RESTART nor NO_RESTART (only in EXTENDED status) E_OS_ACCESS: The caller does not have the right to terminate <Application> (only in EXTENDED status) E_OS_STATE: The state of <Application> does not allow terminating <Application>
Description:	This service terminates the OS-Application to which the calling Task/Category 2 ISR/application specific error hook belongs.	

] ()

[OS493] [If the input parameter <Application> in a call of TerminateApplication() is not valid TerminateApplication() shall return E_OS_ID.] ()

[OS459] [If the <RestartOption> in a call of TerminateApplication() is invalid, TerminateApplication() shall return E_OS_VALUE.] ()

[OS494] [If the input parameter <Application> in a call of TerminateApplication() is valid AND the caller belongs to a non-trusted OS-Application AND the caller does not belong to <Application> TerminateApplication() shall return E_OS_ACCESS.] ()

[OS507] [If the state of <Application> in a call of TerminateApplication() is APPLICATION_TERMINATED TerminateApplication() shall return E_OS_STATE.] ()

[OS508] [If the state of <Application> in a call of TerminateApplication() is APPLICATION_RESTARTING and the caller does not belong to the <Application> then TerminateApplication() shall return E_OS_STATE.] ()

[OS548] [If the state of <Application> in a call of `TerminateApplication()` is `APPLICATION_RESTARTING` AND the caller does belong to the <Application> AND the <RestartOption> is equal `RESTART` then `TerminateApplication()` shall return `E_OS_STATE`.] ()

[OS287] [If the parameters in a call of `TerminateApplication()` are valid and the above criteria are met `TerminateApplication()` shall terminate <Application> (i.e. to kill all tasks, disable the interrupt sources of those ISRs which belong to the OS-Application and free all other OS resources associated with the application) AND shall activate the configured `OsRestartTask` of <Application> if <RestartOption> equals `RESTART`. If the <Application> is restarted, its state is set to `APPLICATION_RESTARTING` otherwise to `APPLICATION_TERMINATED`. If the caller belongs to <Application> `TerminateApplication()` shall not return, otherwise it shall return `E_OK`.] ()

[OS535] [Caveats of `TerminateApplication()`:

- If no applications are configured the implementation shall make sure that this service is not available.
- Tasks and interrupts that are owned by a trusted application can terminate any OS-Application. Tasks and interrupts that are owned by a non-trusted application can only terminate their owning OS-Application.] ()

[OS536] [Availability of `TerminateApplication()`: Available in Scalability Classes 3 and 4.] ()

8.4.20 AllowAccess

[OS501] [

Service name:	AllowAccess	
Syntax:	<pre>StatusType AllowAccess(void)</pre>	
Service ID[hex]:	0x13	
Sync/Async:	Synchronous	
Reentrancy:	Reentrant	
Parameters (in):	None	
Parameters (inout):	None	
Parameters (out):	None	
Return value:	StatusType	E_OK: No errors E_OS_STATE: The OS-Application of the caller is in the wrong state
Description:	This service sets the own state of an OS-Application from <code>APPLICATION_RESTARTING</code> to <code>APPLICATION_ACCESSIBLE</code> .	

] ()

[OS497] [If the state of the OS-Application of the caller of AllowAccess() is not APPLICATION_RESTARTING AllowAccess() shall return E_OS_STATE.] ()

[OS498] [If the state of the OS-Application of the caller of AllowAccess() is APPLICATION_RESTARTING, AllowAccess() shall set the state to APPLICATION_ACCESSIBLE and allow other OS-Applications to access the configured objects of the callers OS-Application.] ()

[OS547] [Availability of AllowAccess(): Available in Scalability Classes 3 and 4.] ()

8.4.21 GetApplicationState

[OS499] [

Service name:	GetApplicationState	
Syntax:	<pre>StatusType GetApplicationState(ApplicationType Application, ApplicationStateRefType Value)</pre>	
Service ID[hex]:	0x14	
Sync/Async:	Synchronous	
Reentrancy:	Reentrant	
Parameters (in):	Application	The OS-Application from which the state is requested
Parameters (inout):	None	
Parameters (out):	Value	The current state of the application
Return value:	StatusType	E_OK: No errors E_OS_ID: <Application> is not valid (only in EXTENDED status)
Description:	This service returns the current state of an OS-Application.	

] ()

[OS495] [If the <Application> in a call of GetApplicationState() is not valid GetApplicationState() shall return E_OS_ID.] ()

[OS496] [If the parameters in a call of GetApplicationState() are valid, GetApplicationState() shall return the state of OS-Application <Application> in <Value>.] ()

[OS537] [Availability of GetApplicationState(): Available in Scalability Classes 3 and 4.] ()

8.4.22 GetNumberOfActivatedCores

[OS672] [

Service name:	GetNumberOfActivatedCores
----------------------	---------------------------

Syntax:	uint32 GetNumberOfActivatedCores(void)
Service ID[hex]:	OS_ServiceID_GetNumberOfActivatedCores
Sync/Async:	Synchronous
Reentrancy:	Reentrant
Parameters (in):	None
Parameters (inout):	None
Parameters (out):	None
Return value:	uint32 Number of cores activated by the StartCore function (see below)
Description:	The function returns the number of cores activated by the StartCore function. This function might be a macro.

] (BSW4080001)

The function `GetNumberOfActivatedCores` shall be callable from within a TASK and an ISR cat 2. Otherwise the behavior is unspecified.

[OS673] [The return value of `GetNumberOfActivatedCores` shall be less or equal to the configured value of “`OsNumberOfCores`”.] (BSW4080001)

8.4.23 GetCoreID

[OS674] [

Service name:	GetCoreID
Syntax:	CoreIdType GetCoreID(void)
Service ID[hex]:	OS_ServiceID_GetCoreID
Sync/Async:	Synchronous
Reentrancy:	Reentrant
Parameters (in):	None
Parameters (inout):	None
Parameters (out):	None
Return value:	CoreIdType The return value is the unique ID of the core.
Description:	The function returns a unique core identifier.

] (BSW4080001)

[OS675] [The function `GetCoreID` shall return the unique logical CoreID of the core on which the function is called. The mapping of physical cores to logical CoreIDs is implementation specific.] (BSW4080001)

8.4.24 StartCore

[OS676] [

Service name:	StartCore
Syntax:	void StartCore(CoreIdType CoreID, StatusType* Status)
Service ID[hex]:	OS_ServiceID_StartCore

Sync/Async:	Synchronous	
Reentrancy:	Non Reentrant	
Parameters (in):	CoreID	Core identifier
Parameters (inout):	None	
Parameters (out):	Status	Return value of the function in extended status: E_OK: No Error E_OS_ID: Core ID is invalid. E_OS_ACCESS: The function was called after starting the OS. E_OS_STATE: The Core is already activated. Return value of the function in standard status E_OK: No Error
Return value:	None	
Description:	It is not supported to call this function after StartOS(). The function starts the core specified by the parameter CoreID. The OUT parameter allows the caller to check whether the operation was successful or not. If a core is started by means of this function StartOS shall be called on the core.	

] (BSW4080006, BSW4080026, BSW4080027)

[OS677] [The function StartCore shall start one core that shall run under the control of the AUTOSAR OS.] (BSW4080006, BSW4080026, BSW4080027)

[OS678] [Calls to the StartCore function after StartOS() shall return with E_OS_ACCESS and the core shall not be started.] (BSW4080006, BSW4080026, BSW4080027)

[OS679] [If the parameter CoreIDs refers to a core that was already started by the function StartCore the related core is ignored and E_OS_STATE shall be returned.] (BSW4080006, BSW4080026, BSW4080027)

[OS680] [If the parameter CoreID refers to a core that was already started by the function StartNonAutosarCore the related core is ignored and E_OS_STATE shall be returned.] (BSW4080006, BSW4080026, BSW4080027)

[OS681] [There is no call to the ErrorHandler() if an error occurs during StartCore();] (BSW4080006, BSW4080026, BSW4080027)

8.4.25 StartNonAutosarCore

[OS682] [

Service name:	StartNonAutosarCore
Syntax:	<pre>void StartNonAutosarCore(CoreIdType CoreID, StatusType* Status)</pre>
Service ID[hex]:	OS_ServiceID_StartNonAutosarCore
Sync/Async:	Synchronous

Reentrancy:	Non Reentrant	
Parameters (in):	CoreID	Core identifier
Parameters (inout):	None	
Parameters (out):	Status	Return value of the function in standard status: E_OK: No Error E_OS_ID: Core ID is invalid. E_OS_STATE: The Core is already activated. Return value of the function in extended status E_OK: No Error
Return value:	None	
Description:	The function starts the core specified by the parameter CoreID. It is allowed to call this function after StartOS(). The OUT parameter allows the caller to check whether the operation was successful or not. It is not allowed to call StartOS on cores activated by StartNonAutosarCore. Otherwise the behaviour is unspecified.	

] (BSW4080006, BSW4080026, BSW4080027)

[OS683] [The function `StartNonAutosarCore` shall start a core that is not controlled by the AUTOSAR OS.] (BSW4080006, BSW4080026, BSW4080027)

[OS684] [If the parameter `CoreID` refers to a core that was already started by the function `StartNonAutosarCore` has no effect and sets "Status" to `E_OS_STATE`.

] (BSW4080006, BSW4080026, BSW4080027)

[OS685] [If the parameter `CoreID` refers to an unknown core the function `StartNonAutosarCore` has no effect and sets "Status" to `E_OS_ID`.] (BSW4080006, BSW4080026, BSW4080027)

8.4.26 GetSpinlock

[OS686] [

Service name:	GetSpinlock	
Syntax:	<pre>StatusType GetSpinlock(SpinlockIdType SpinlockId)</pre>	
Service ID[hex]:	OS_ServiceID_GetSpinlock	
Sync/Async:	Synchronous	
Reentrancy:	Reentrant	
Parameters (in):	SpinlockId	The value refers to the spinlock instance that shall be locked.
Parameters (inout):	None	
Parameters (out):	None	
Return value:	StatusType	E_OK - In standard and extended status : No Error E_OS_ID - In extended status: The SpinlockId is invalid E_OS_INTERFERENCE_DEADLOCK - In extended status: A TASK tries to occupy the spinlock while the lock is already occupied by a TASK on the same core. This would cause a deadlock. E_OS_NESTING_DEADLOCK - In extended status: A TASK tries to occupy the spinlock while holding a different spinlock in a way that may cause a deadlock.

	E_OS_ACCESS - In extended status: The spinlock cannot be accessed.
Description:	GetSpinlock tries to occupy a spin-lock variable. If the function returns, either the lock is successfully taken or an error has occurred. The spinlock mechanism is an active polling mechanism. The function does not cause a de-scheduling.

] (BSW4080021)

[OS687] [The function `GetSpinlock` shall occupy a spinlock. If the spinlock is already occupied the function shall busy wait until the spinlock becomes available.] (BSW4080021)

[OS688] [The function `GetSpinlock` shall return `E_OK` if no error was detected. The spinlock is now occupied by the calling TASK/ISR2 on the calling core.] (BSW4080021)

[OS689] [The function `GetSpinlock` shall return `E_OS_ID` if the parameter `SpinlockID` refers to a spinlock that does not exist.] (BSW4080021)

[OS690] [The function `GetSpinlock` shall return `E_OS_INTERFERENCE_DEADLOCK` if the spinlock referred by the parameter `SpinlockID` is already occupied by a TASK/ISR2 on the same core.] (BSW4080021)

[OS691] [The function `GetSpinlock` shall return `E_OS_NESTING_DEADLOCK` if the sequence by which multiple spinlocks are occupied at the same time do not comply with the configured order.] (BSW4080021)

[OS692] [The function `GetSpinlock` shall return `E_OS_ACCESS` if the accessing OS-Application was not listed in the configuration (`OsSpinlock`).] (BSW4080021)

[OS693] [It shall be allowed to call the function `GetSpinlock` while interrupts are disabled.] (BSW4080021)

[OS694] [It shall be allowed to call the function `GetSpinlock` while a RESOURCE is occupied.] (BSW4080021)

8.4.27 ReleaseSpinlock

[OS695] [

Service name:	ReleaseSpinlock
Syntax:	StatusType ReleaseSpinlock(SpinlockIdType SpinlockId)
Service ID[hex]:	OS_ServiceID_ReleaseSpinlock
Sync/Async:	Synchronous
Reentrancy:	Reentrant

Parameters (in):	SpinlockId	The value refers to the spinlock instance that shall be locked.
Parameters (inout):	None	
Parameters (out):	None	
Return value:	StatusType	E_OK - In standard and extended status: No Error E_OS_ID - In extended status: The SpinlockId is invalid. E_OS_STATE - In extended status: The Spinlock is not occupied by the TASK E_OS_ACCESS - In extended status: The Spinlock cannot be accessed. E_OS_NOFUNC - In extended status: Attempt to release a spinlock while another spinlock has to be released before.
Description:	ReleaseSpinlock releases a spinlock variable that was occupied before. Before terminating a TASK all spinlock variables that have been occupied with GetSpinlock() shall be released. Before calling WaitEVENT all Spinlocks shall be released.	

] (BSW4080021)

[OS696] [The function `ReleaseSpinlock` shall release a spinlock that has been occupied by the same (calling) TASK.] (BSW4080021)

[OS697] [The function `ReleaseSpinlock` shall return `E_OK` if no error was detected. The spinlock is now free and can be occupied by the same or other TASKS.

] (BSW4080021)

[OS698] [The function `ReleaseSpinlock` shall return `E_OS_ID` if the parameter `SpinlockID` refers to a spinlock that does not exist.] (BSW4080021)

[OS699] [The function `ReleaseSpinlock` shall return `E_OS_STATE` if the parameter `SpinlockID` refers to a spinlock that is not occupied by the calling TASK.

] (BSW4080021)

[OS700] [The function `ReleaseSpinlock` shall return `E_OS_ACCESS` if the TASK has no access to the spinlock referred by the parameter `SpinlockID`] (BSW4080021)

[OS701] [The function `ReleaseSpinlock` shall return `E_OS_NOFUNC` if the TASK tries to release a spinlock while another spinlock has to be released before. No functionality shall be performed.] (BSW4080021)

[OS702] [Spinlocks and RESOURCES can only be locked and unlocked in strict LIFO order. Otherwise `E_OS_RESOURCE` shall be returned.] (BSW4080021)

8.4.28 TryToGetSpinlock

[OS703] [

Service name:	TryToGetSpinlock
----------------------	------------------

Syntax:	<pre>StatusType TryToGetSpinlock(SpinlockIdType SpinlockId, TryToGetSpinlockType* Success)</pre>	
Service ID[hex]:	OS_ServiceID_TryToGetSpinlock	
Sync/Async:	Synchronous	
Reentrancy:	Reentrant	
Parameters (in):	SpinlockId	The value refers to the spinlock instance that shall be locked.
Parameters (inout):	None	
Parameters (out):	Success	Returns if the lock has been occupied or not
Return value:	StatusType	<p>E_OK - In standard and extended status: No Error</p> <p>E_OS_ID - In extended status: The SpinlockId is invalid.</p> <p>E_OS_INTERFERENCE_DEADLOCK - In extended status: A TASK tries to occupy the spinlock while the lock is already occupied by a TASK on the same core. This would cause a deadlock.</p> <p>E_OS_NESTING_DEADLOCK - In extended status: A TASK tries to occupy a spinlock while holding a different spinlock in a way that may cause a deadlock.</p> <p>E_OS_ACCESS - In extended status: The spinlock cannot be accessed.</p>
Description:	TryToGetSpinlock has the same functionality as GetSpinlock with the difference that if the spinlock is already occupied by a TASK on a different core the function sets the OUT parameter "Success" and returns with E_OK.	

] (BSW4080021)

[OS704] [The function TryToGetSpinlock shall atomically test the availability of the spinlock and if available occupy it. The result of success is returned.] (BSW4080021)

[OS705] [The function TryToGetSpinlock shall set the OUT parameter "Success" to TRYTOGETSPINLOCK_SUCCESS if the spinlock was successfully occupied, and TRYTOGETSPINLOCK_NOSUCCESS if not. In both cases E_OK shall be returned.] (BSW4080021)

[OS706] [If the function TryToGetSpinlock does not return E_OK, the OUT parameter "Success" shall be undefined.] (BSW4080021)

[OS707] [The function TryToGetSpinlock shall return E_OS_ID if the parameter SpinlockID refers to a spinlock that does not exist.] (BSW4080021)

[OS708] [The function TryToGetSpinlock shall return E_OS_INTERFERENCE_DEADLOCK if the spinlock referred by the parameter SpinlockID is already occupied by a TASK on the same core.] (BSW4080021)

[OS709] [The function TryToGetSpinlock shall return E_OS_NESTING_DEADLOCK if a TASK tries to occupy a spinlock while holding a different spinlock in a way that may cause a deadlock.] (BSW4080021)

[OS710] [The function `TryToGetSpinlock` shall return `E_OS_ACCESS` if the `TASK` has no access to the spinlock referred by the parameter `SpinlockID`] (BSW4080021)

[OS711] [It shall be allowed to call the function `TryToGetSpinlock` while interrupts are disabled.] (BSW4080021)

[OS712] [It shall be allowed to call the function `TryToGetSpinlock` while a `RESOURCE` is occupied.] (BSW4080021)

8.4.29 ShutdownAllCores

[OS713] [

Service name:	ShutdownAllCores
Syntax:	<code>void ShutdownAllCores(StatusType Error)</code>
Service ID[hex]:	OS_ServiceID_ShutdownAllCores
Sync/Async:	Synchronous
Reentrancy:	Reentrant
Parameters (in):	Error <Error> needs to be a valid error code supported by the AUTOSAR OS.
Parameters (inout):	None
Parameters (out):	None
Return value:	None
Description:	After this service the OS on all AUTOSAR cores is shut down. Allowed at <code>TASK</code> level and <code>ISR</code> level and also internally by the OS. The function will never return. The function will force other cores into a shutdown.

] (BSW4080007)

[OS714] [A Synchronized shutdown shall be triggered by the API function `ShutdownAllCores`.] (BSW4080007)

[OS715] [`ShutdownAllCores` shall not return.] (BSW4080007)

[OS716] [If `ShutdownAllCores` is called from non trusted code the call shall be ignored.] (BSW4080007)

8.5 IOC

8.5.1 Imported types

In this chapter all types included from the following files are listed:

<i>Module</i>	<i>Imported Type</i>
GENERIC TYPES	...
	<Data1>
	<Data2>
	<Data>
Std_Types	Std_ReturnType

8.5.2 Type definitions

None

8.5.3 Constants

Name	Communication	Type	Errorname / Value	Annotation
IOC_E_OK	All, SND/RCV	Std_ReturnType	RTE_E_OK / 0	No error occurred
IOC_E_NOK	All SND/RCV	Std_ReturnType	RTE_E_NOK / 1	Error occurred. Shall be used to identify error cases without error specification.
IOC_E_LIMIT	Queued SND	Std_ReturnType	RTE_E_LIMIT / 130	In case of "event" (queued) semantic, the internal buffer within the IOC communication service is full (Case: Receiver slower than sender). This error produces additionally an Overlaid Error on the receiver side at the next data reception.
IOC_E_LOST_DATA	Queued RCV	Std_ReturnType	Overlaid Error RTE_E_LOST_DATA / 64	In case of "event" (queued) semantic, this Overlaid Error indicates that the IOC service refuses an locSend request due to internal buffer overflow.
IOC_E_NO_DATA	Queued RCV	Std_ReturnType	RTE_E_NO_DATA / 131	In case of "event" (queued) semantic, no data is available for reception.

8.5.4 Function definitions

8.5.4.1 locSend/locWrite

The `IocWrite` API call is generated for "data" (unqueued) semantics and the `IocSend` API call is generated for "events" (queued) semantics.

[OS718] [

Service name:	locSend_<locId>[_<SenderId>]	
Syntax:	Std_ReturnType IocSend_<IocId>[_<SenderId>](<Data> IN)	
Service ID[hex]:	IOCServId_IOC_Send	
Sync/Async:	Asynchronous	
Reentrancy:	This function is generated individually for each sender. The individual function is not reentrant (if called from different runnable entities that belong to the same sender), but different functions can be called in parallel.	
Parameters (in):	IN	Data value to be sent over a communication identified by the <locId>. The parameter will be passed by value for primitive data elements and by reference for all other types. Example: Std_ReturnType locSend_RTE_25 (const uint32 UI_Value); Std_ReturnType locSend_RTE_42 (const TASKParams3 *pStr_Value);
Parameters (inout):	None	
Parameters (out):	None	
Return value:	Std_ReturnType	IOC_E_OK: The data has been passed successfully to the communication service. IOC_E_LIMIT: IOC internal communication buffer is full (Case: Receiver is slower than sender). This error produces an IOC_E_LOST_DATA Overlaid Error on the receiver side at the next data reception.
Description:	Performs an "explicit" sender-receiver transmission of data elements with "event" semantic for a unidirectional 1:1 or N:1 communication between OS-Applications located on the same or on different cores. <locId> is a unique identifier that references a unidirectional 1:1 or N:1 communication. <SenderId> is used only in N:1 communication. Together with <locId>, it uniquely identifies the sender. It is separated from <locId> with an underscore. In case of 1:1 communication, it shall be omitted.	

Service name:	locWrite_<locId>[_<SenderId>]	
Syntax:	Std_ReturnType IocWrite_<IocId>[_<SenderId>](<Data> IN)	

Service ID[hex]:	IOCServiceld_IOC_Write	
Sync/Async:	Asynchronous	
Reentrancy:	This function is generated individually for each sender. The individual function is not reentrant (if called from different runnable entities that belong to the same sender), but different functions can be called in parallel.	
Parameters (in):	IN	Data value to be sent over a communication identified by the <locId>. The parameter will be passed by value for primitive data elements and by reference for all other types. Example: Std_ReturnType locWrite_RTE_25 (const uint32 UI_Value); Std_ReturnType locWrite_RTE_42 (const TASKParams3 *pStr_Value);
Parameters (inout):	None	
Parameters (out):	None	
Return value:	Std_ReturnType	IOC_E_OK: The data has been passed successfully to the communication service.
Description:	Performs an "explicit" sender-receiver transmission of data elements with "data" semantic for a unidirectional 1:1 or N:1 communication between OS-Applications located on the same or on different cores. <locId> is a unique identifier that references a unidirectional 1:1 or N:1 communication. <SenderId> is used only in N:1 communication. Together with <locId>, it uniquely identifies the sender. It is separated from <locId> with an underscore. In case of 1:1 communication, it shall be omitted.	

] (BSW4080020)

General:

[OS719] [IocSend/IocWrite is asynchronous in that way it shall not have to wait for the reception of the data on the receiving side to return from execution.] (BSW4080020)

[OS720] [The IocSend/IocWrite function shall not return until the data given in parameter have been completely physically sent over the communication medium.

For example in case of communication over shared RAM, an IocSend/IocWrite shall return when all data have been copied in the target shared RAM.] (BSW4080020)

[OS721] [In case of "event" (queued) semantic, the IocSend function shall guarantee the order of delivery. In case of senders from different cores, the order in which messages are received will be determined by the implementation.] (BSW4080020)

[OS722] [The IocSend/IocWrite function shall support mechanism to guarantee data-Integrity during transmission.

The IocSend/IocWrite function shall solve the crossing of the protection boundaries of OS-Applications. It has to be generated in case of intra-core and inter-core communication.] (BSW4080020)

Parameters:

[OS723] [The IN <Data> parameter of the `IocSend/IocWrite` function shall be passed by value for primitive data types and by reference for all other types.] (BSW4080020)

[OS724] [For data passed by reference, the `IocSend/IocWrite` function shall guarantee upon return that the parameter reference is safe for re-use.] (BSW4080020)

Returned values:

[OS725] [The `IocSend/IocWrite` function shall return `IOC_E_OK` if the data was passed successfully to the communication service.] (BSW4080020)

[OS726] [In case of “event” semantic the `IocSend` function shall return `IOC_E_LIMIT` if an IOC internal transmission buffer became full (Case: Receiver is slower than sender or/and configured internal IOC buffer size is too small). If this error occurs the IOC internal buffer could not be filled with the parameter. In that case this error shall produce an `IOC_E_LOST_DATA` Overlayed Error on the receiver side at the next data reception (s. OS745).] (BSW4080020)

Internal structures:

[OS727] [In case of “event” semantic the IOC shall configure its internal transmission buffer size with the value of the attribute `OsIocBufferLength`.] (BSW4080020)

8.5.4.2 `IocSendGroup/IocWriteGroup`

The `IocWriteGroup` API call is generated for "data" (unqueued) semantics and the `IocSendGroup` API call is generated for "events" (queued) semantics.

[OS728] [

Service name:	<code>IocSendGroup_<IocId></code>	
Syntax:	<pre>Std_ReturnType IocSendGroup_<IocId>(<Data1> IN1, <Data2> IN2, ...)</pre>	
Service ID[hex]:	<code>IOCServiceld_IOC_SendGroup</code>	
Sync/Async:	Asynchronous	
Reentrancy:	This function is generated individually for each sender. The individual function is not reentrant (if called from different runnable entities that belong to the same sender), but different functions can be called in parallel.	
Parameters (in):	IN1	List of parameters with data values to be sent over a communication identified by the <IocId>. The parameters will be passed by value for simple data elements and by reference for all other types.

		Example: Std_ReturnType locSendGroup_RTE_G1 (const uint32 UI_Value1, const uint16 Value2, const uint8 Value3, const uint16 Value4);
	IN2	--
	--	--
Parameters (inout):	None	
Parameters (out):	None	
Return value:	Std_ReturnType	IOC_E_OK: The data has been passed successfully to the communication service. IOC_E_LIMIT: IOC internal communication buffer is full (Case: Receiver is slower than sender). This error produces an IOC_E_LOST_DATA Overlaid Error on the receiver side at the next data reception.
Description:	<p>Performs an "explicit" sender-receiver transmission of data elements with "event" semantic for a unidirectional 1:1 communication between OS-Applications located on the same or on different cores.</p> <p>This API involves a group of data elements which values are specified in parameter.</p> <p><locId> is a unique identifier that references a unidirectional 1:1 communication involving many data elements.</p>	

Service name:	locWriteGroup_<locId>	
Syntax:	Std_ReturnType IocWriteGroup_<IocId>(<Data1> IN1, <Data2> IN2, ...)	
Service ID[hex]:	IOCServicId_IOC_WriteGroup	
Sync/Async:	Asynchronous	
Reentrancy:	This function is generated individually for each sender. The individual function is not reentrant (if called from different runnable entities that belong to the same sender), but different functions can be called in parallel.	
Parameters (in):	IN1	List of parameters with data values to be sent over a communication identified by the <locId>. The parameters will be passed by value for simple data elements and by reference for all other types. Example: Std_ReturnType locWriteGroup_RTE_G1 (const uint32 UI_Value1, const uint16 Value2, const uint8 Value3, const uint16 Value4);
	IN2	--
	--	--
Parameters (inout):	None	
Parameters (out):	None	
Return value:	Std_ReturnType	IOC_E_OK: The data has been passed successfully to the communication service.
Description:	Performs an "explicit" sender-receiver transmission of data elements with "data"	

	<p>semantic for a unidirectional 1:1 communication between OS-Applications located on the same or on different cores.</p> <p>This API involves a group of data elements which values are specified in parameter.</p> <p><loclId> is a unique identifier that references a unidirectional 1:1 communication involving many data elements.</p>
--	--

] (BSW4080020)

General:

[OS729] [IocSendGroup/IocWriteGroup is asynchronous in that way it shall not have to wait for the reception of the data on the receiving side to return from execution.] (BSW4080020)

[OS730] [The IocSendGroup/IocWriteGroup function shall not return until the data given in parameter have been completely physically sent over the communication medium. For example in case of communication over shared RAM, an IocSendGroup/IocWriteGroup shall return when all data have been copied in the target shared RAM.] (BSW4080020)

[OS731] [In case of “event” semantic, the IocSendGroup function shall guarantee the order of delivery.] (BSW4080020)

[OS732] [The IocSendGroup/IocWriteGroup function shall support mechanisms to guarantee data-Integrity during transmission.

The IocSendGroup/IocWriteGroup function shall solve the crossing of the protection boundaries of OS-Applications. It has to be generated in case of intra-core and inter-core communication.] (BSW4080020)

Parameters:

[OS733] [A parameter IN <Data> may be passed by value for simple data elements and by reference for all other types.] (BSW4080020)

[OS734] [For data passed by reference, the IocSendGroup/IocWriteGroup function shall guarantee upon return that the parameter reference is safe for re-use.] (BSW4080020)

Returned values:

[OS735] [The IocSendGroup/IocWriteGroup function shall return IOC_E_OK if the data was passed successfully to the communication service.] (BSW4080020)

[OS736] [In case of “event” semantic the IocSendGroup function shall return IOC_E_LIMIT if an IOC internal transmission buffer got full (Case: Receiver is slower than sender or/and configured internal IOC buffer size is too small).

If this error occurs the IOC Internal buffer could not be filled with the parameter. In that case this error produces an IOC_E_LOST_DATA Overlaid Error on the receiver side at the next data reception.] (BSW4080020)

Internal structures:

[OS737] [In case of "event" semantic the IOC shall configure its internal transmission buffer size with the value of the attribute `OsIocBufferLength`.] (BSW4080020)

8.5.4.3 IocReceive/IocRead

The `IocRead` API call is generated for "data" and the `IocReceive` API call is generated for "events".

[OS738] [

Service name:	IocReceive_<IocId>	
Syntax:	Std_ReturnType IocReceive_<IocId>(<Data> OUT)	
Service ID[hex]:	IOCServiceld_IOC_Receive	
Sync/Async:	Synchronous	
Reentrancy:	This function is generated individually for each receiver. The individual function is not reentrant (if called from different runnable entities that belong to the same receiver), but different functions can be called in parallel.	
Parameters (in):	None	
Parameters (inout):	None	
Parameters (out):	OUT	Data reference to be filled with the received data element.
Return value:	Std_ReturnType	IOC_E_OK: Data was received successfully IOC_E_NO_DATA: No data is available for reception. IOC_E_LOST_DATA: This Overlaid Error indicates that the IOC communication service refused an IOCSend request from sender due to an internal buffer overflow. There is no error in the data returned in parameter.
Description:	Performs an "explicit" sender-receiver reception of data elements with "event" semantic for a unidirectional communication between OS-Applications located on the same or on different cores.. <IocId> is a unique identifier that references a unidirectional 1:1 or N:1 communication.	

Service name:	IocRead_<IocId>	
Syntax:	Std_ReturnType IocRead_<IocId>(<Data> OUT)	
Service ID[hex]:	IOCServiceld_IOC_Read	
Sync/Async:	Synchronous	
Reentrancy:	This function is generated individually for each receiver. The individual function is not reentrant (if called from different runnable entities that belong to the same receiver), but different functions can be called in parallel.	

Parameters (in):	None	
Parameters (inout):	None	
Parameters (out):	OUT	Data reference to be filled with the received data element.
Return value:	Std_ReturnType	IOC_E_OK: Data was received successfully
Description:	<p>Performs an "explicit" sender-receiver reception of data elements with "data" semantic for a unidirectional communication between OS-Applications located on the same or on different cores.</p> <p><locId> is a unique identifier that references a unidirectional 1:1 or N:1 communication.</p>	

] (BSW4080020)

General:

[OS739] [A successful call to the `IocReceive/IocRead` function indicates that data has been received successfully in the OUT <Data> given in parameter.

The `IocReceive/IocRead` function has to be generated in case of intra-core and inter-core communication.] (BSW4080020)

[OS740] [If the `OsIocReceiverPullCB` attribute is defined with a callback function name, the IOC shall call this function on the receiving core for each data transmission.] (BSW4080020)

Parameters:

[OS741] [In case of "data" semantic the `IocRead` function shall always be able to deliver the last available datum. In case of senders from different cores, the precision of the order might be limited by the hardware and implementation.] (BSW4080020)

[OS742] [The `IocReceive/IocRead` function shall guarantee upon returning from execution that the reference given in parameter is safe for use.] (BSW4080020)

Returned values:

[OS743] [The `IocReceive/IocRead` function shall return `IOC_E_OK` if the data was received successfully in the OUT <Data> parameter.] (BSW4080020)

[OS744] [In case of "event" semantic and if no data is available the function `IocReceive` shall return `IOC_E_NO_DATA`.] (BSW4080020)

[OS745] [In case of "event" semantic an `IOC_E_LOST_DATA` Overlayed Error shall be returned by the `IocReceive` function if the IOC communication service refused an `IocSend` request from sender due to an internal buffer overflow. There is no error in the data returned in parameter.] (BSW4080020)

8.5.4.4 IocReceiveGroup/IocReadGroup

The IocReadGroup API call is generated for "data" and the IocReceiveGroup API call is generated for "events".

[OS746] [

Service name:	IocReceiveGroup_<IocId>	
Syntax:	<pre>Std_ReturnType IocReceiveGroup_<IocId>(<Data1> OUT1, <Data2> OUT2, ...)</pre>	
Service ID[hex]:	IOCServId_IOC_ReceiveGroup	
Sync/Async:	Synchronous	
Reentrancy:	This function is generated individually for each receiver. The individual function is not reentrant (if called from different runnable entities that belong to the same receiver), but different functions can be called in parallel.	
Parameters (in):	None	
Parameters (inout):	None	
Parameters (out):	OUT1	List of data references to be filled with the received data elements. The specified order of the parameter shall match to the specified order in the corresponding send function.
	OUT2	--
	--	--
Return value:	Std_ReturnType	IOC_E_OK: Data was received successfully IOC_E_NO_DATA: No data is available for reception. IOC_E_LOST_DATA: This Overlaid Error indicates that the IOC communication service refused an IOCSend request from sender due to an internal buffer overflow. There is no error in the data returned in parameter.
Description:	Performs an "explicit" sender-receiver transmission of data elements with "event" semantic for a unidirectional 1:1 communication between OS-Applications located on the same or on different cores. This API involves a group of data elements which values are specified in parameter. <IocId> is a unique identifier that references a unidirectional 1:1 communication involving many data elements.	

Service name:	IocReadGroup_<IocId>	
Syntax:	<pre>Std_ReturnType IocReadGroup_<IocId>(<Data1> OUT1, <Data2> OUT2, ...)</pre>	
Service ID[hex]:	IOCServId_IOC_ReadGroup	
Sync/Async:	Synchronous	
Reentrancy:	This function is generated individually for each receiver. The individual function is not reentrant (if called from different runnable entities that belong to the same	

	receiver), but different functions can be called in parallel.	
Parameters (in):	None	
Parameters (inout):	None	
Parameters (out):	OUT1	List of data references to be filled with the received data elements. The specified order of the parameter shall match to the specified order in the corresponding send function.
	OUT2	--
	--	--
Return value:	Std_ReturnType	IOC_E_OK: Data was received successfully
Description:	<p>Performs an "explicit" sender-receiver transmission of data elements with a "data" semantic for a unidirectional 1:1 communication between OS-Applications located on the same or on different cores.</p> <p>This API involves a group of data elements which values are specified in parameter.</p> <p><locl> is a unique identifier that references a unidirectional 1:1 communication involving many data elements.</p>	

] (BSW4080020)

General:

[OS747] [A successful call to the `IocReceiveGroup/IocReadGroup` function indicates that data has been received successfully in the given parameters.

The `IocReceiveGroup/IocReadGroup` function has to be generated in case of intra-core and inter-core communication.] (BSW4080020)

[OS748] [If the `OsIocReceiverPullCB` attribute is defined with a callback function name, the IOC shall call this function on the receiving core for each data transmission.] (BSW4080020)

Parameters:

[OS749] [In case of "data" semantic the `IocReadGroup` function shall always be able to deliver the last available datum.] (BSW4080020)

[OS750] [The `IocReceiveGroup/IocReadGroup` function shall guarantee upon returning from execution that the references given in parameters are safe for use.] (BSW4080020)

Returned values:

[OS751] [The `IocReceiveGroup/IocReadGroup` function shall return `IOC_E_OK` if the data was received successfully in the list of references given in parameter.] (BSW4080020)

[OS752] [In case of "event" semantic and if no data is available the function `IocReceiveGroup` shall return `IOC_E_NO_DATA`.] (BSW4080020)

[OS753] [In case of “event” semantic an IOC_E_LOST_DATA Overlaid Error shall be returned by the `IocReceiveGroup` function if the IOC communication service refused an `IocSendGroup` request from sender due to an internal buffer overflow. There is no error in the data returned in parameter.] (BSW4080020)

8.5.4.5 IocEmptyQueue

[OS754] [

Service name:	IocEmptyQueue_<IocId>	
Syntax:	Std_ReturnType IocEmptyQueue_<IocId>(void)	
Service ID[hex]:	IOCServiceld_IOC_EmptyQueue	
Sync/Async:	Synchronous	
Reentrancy:	Non reentrant	
Parameters (in):	None	
Parameters (inout):	None	
Parameters (out):	None	
Return value:	Std_ReturnType	IOC_E_OK: Content of the queue was successfully deleted
Description:	In case of queued communication identified by the <IocId> in the function name, the content of the IOC internal communication queue shall be deleted.	

] (BSW4080020)

General:

[OS755] [The function `IocEmptyQueue_<IocId>` shall be present for all IOC elements with queued semantics.] (BSW4080020)

[OS756] [The function `IocEmptyQueue_<IocId>` shall delete all contents from the associated data queue.

The `IocEmptyQueue` should be generated in a more efficient way than an iterative call to an `IocReceive` function.] (BSW4080020)

8.6 Expected Interfaces

In this chapter all interfaces required from other modules are listed.

8.6.1 Mandatory Interfaces

There are no mandatory interfaces for the IOC.

8.6.2 Optional Interfaces

8.6.2.1 ReceiverPullCB

[OS757] [

Service name:	<ReceiverPullCB>
Syntax:	void <ReceiverPullCB>(void)
Service ID[hex]:	IOCServId_IOC_ReceiverPullCB
Sync/Async:	Synchronous
Reentrancy:	Reentrant
Parameters (in):	None
Parameters (inout):	None
Parameters (out):	None
Return value:	None
Description:	This callback function can be configured for the receiver of a communication. If configured, IOC calls this callback on the receiving core for each data reception. <ReceiverPullCB> is the callback function name configured by the receiver in the OsIocReceiverPullCB attribute to be called on data reception."

] (BSW4080020)

[OS758] [The <ReceiverPullCB> function name shall be defined within a configuration file for each IOC communication in the OsIocReceiverPullCB attribute.] (BSW4080020)

[OS759] [The name of the callback shall be unique over the micro controller. For this purpose the following example can be considered as orientation for the IOC user:

Example: Rte_IocReceiveCB_<IocId>] (BSW4080020)

[OS760] [The <ReceiverPullCB> function on the receiver side shall have the same calling rights as a category 2 ISR.] (BSW4080020)

[OS761] [This notification mechanism shall be supported for both queued and unqueued communication semantic.] (BSW4080020)

The owner of the <ReceiverPullCB> function shall pay attention that the execution time of the function shall not last too long. It shall be possible to call this function from an IOC-ISR.

8.7 Hook functions

Hook functions are called by the operating system if specific conditions are met. They are provided by the user. Besides the ProtectionHook below, the hooks from [17] and/or extensions from 7.12 may be called by the OS.

8.7.1 Protection Hook

[OS538] [
Service name:	ProtectionHook	
Syntax:	ProtectionReturnType ProtectionHook(StatusType Fatalerror)	
Service ID[hex]:	0x00	
Sync/Async:	Synchronous	
Reentrancy:	Reentrant	
Parameters (in):	Fatalerror	The error which caused the call to the protection hook
Parameters (inout):	None	
Parameters (out):	None	
Return value:	ProtectionReturnType	PRO_IGNORE PRO_TERMINATETASKISR PRO_TERMINATEAPPL PRO_TERMINATEAPPL_RESTART PRO_SHUTDOWN The return value defines the action the OS shall take after the protection hook.
Description:	The protection hook is always called if a serious error occurs. E.g. exceeding the worst case execution time or violating against the memory protection.	

] ()

Depending on the return value the Operating System module will either:

- forcibly terminate the Task/Category 2 ISR which causes the problem OR
- forcibly terminate the OS-Application the Task/Category 2 ISR belong (optional with restart) OR
- shutdown the system OR
- do nothing

(see 7.8.2)

[OS308] [If ProtectionHook() returns an invalid value, the Operating System module shall take the same action as if no protection hook is configured.] ()

[OS542] [Availability of ProtectionHook(): Available in Scalability Classes 2, 3 and 4.] ()

8.7.2 Application specific StartupHook

[OS539] [

Service name:	StartupHook_<App>
Syntax:	void StartupHook_<App>(void)
Service ID[hex]:	0x00
Sync/Async:	Synchronous
Reentrancy:	Reentrant
Parameters (in):	None
Parameters (inout):	None
Parameters (out):	None
Return value:	None
Description:	The application specific startup hook is called during the start of the OS (after the user has started the OS via StartOS()).

The application specific StartupHook is always called after the standard StartupHook() (see [OS236](#)). If more than one OS-Application is configured which use startup hooks, the order of calls to the startup hooks of the different OS-Applications is not defined.] ()

[OS543] [Availability of StartupHook_<App>(): Available in Scalability Classes 3 and 4.] ()

8.7.3 Application specific ErrorHandler

[OS540] [

Service name:	ErrorHook_<App>
Syntax:	void ErrorHandler_<App>(StatusType Error)
Service ID[hex]:	0x00
Sync/Async:	Synchronous
Reentrancy:	Reentrant
Parameters (in):	Error The error which caused the call to the error hook
Parameters (inout):	None
Parameters (out):	None
Return value:	None
Description:	The application specific error hook is called whenever a Task or Category 2 ISR which belongs to the OS-Application causes an error.

If the general ErrorHandler() is configured, the general ErrorHandler() is called before the application specific error hook is called (see [OS246](#)).] ()

[OS544] [Availability of ErrorHandler_<App>(): Available in Scalability Classes 3 and 4.] ()

8.7.4 Application specific ShutdownHook

[OS541] [

Service name:	ShutdownHook_<App>
Syntax:	void ShutdownHook_<App>(StatusType Fatalerror)
Service ID[hex]:	0x00
Sync/Async:	Synchronous
Reentrancy:	Reentrant
Parameters (in):	Fatalerror The error which caused the action to shut down the operating system.
Parameters (inout):	None
Parameters (out):	None
Return value:	None
Description:	The application specific shutdown hook is called whenever the system starts the shut down of itself.

If the general ShutdownHook() is configured, the general ShutdownHook() is called after all application specific shutdown hook(s) are called (see [OS237](#)). If more OS-Applications with an application specific shutdown hook exist the order of calls to these application specific shutdown hooks is not defined.] ()

[OS545] [Availability of ShutdownHook_<App>(): Available in Scalability Classes 3 and 4.] ()

9 Sequence diagrams

9.1 Sequence chart for calling trusted functions

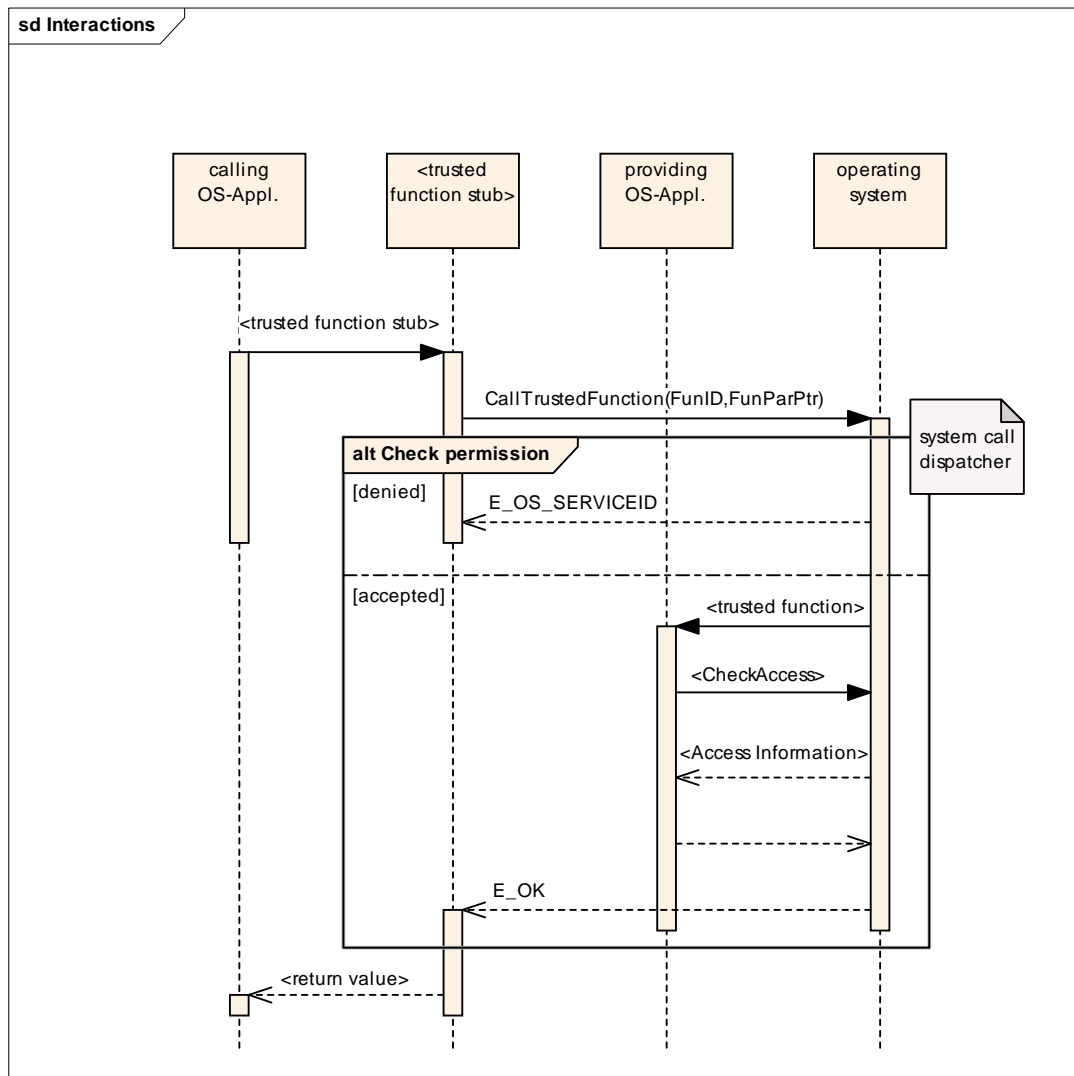


Figure 9.1: System Call sequence chart

The above sequence describes a call to the `CallTrustedFunction` service. It starts with a user who calls a service which requires itself a call to a trusted function. The service then packs the argument for the trusted function into a structure and calls `CallTrustedFunction` with the ID and the pointer as arguments. Afterwards the OS checks if the access to the requested service is valid. If no access is granted `E_OS_SERVICEID` is returned. Otherwise the trusted service itself is called and the function checks the arguments for access right, etc.

9.2 Sequence chart for usage of ErrorHook

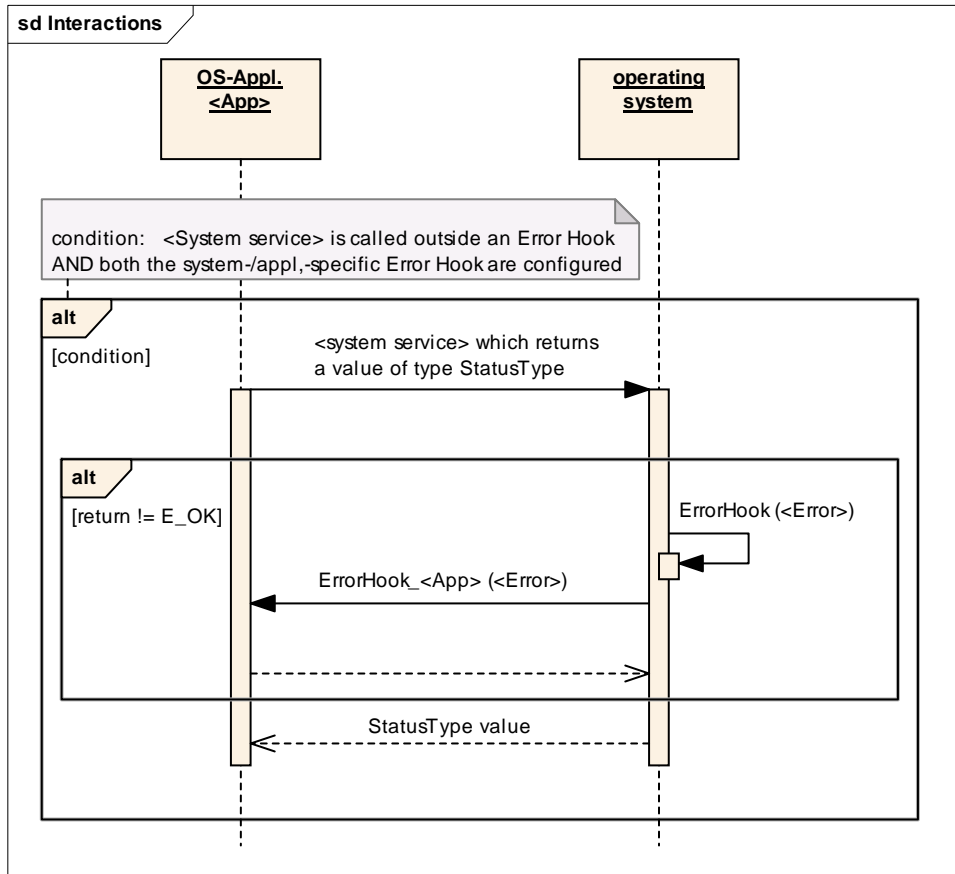


Figure 9.2: Error Hook sequence chart

The above sequence chart shows the sequence of error hook calls in case a service does not return with `E_OK`. Note that in this case the general error hook and the OS-Application specific error hook are called.

9.3 Sequence chart for ProtectionHook

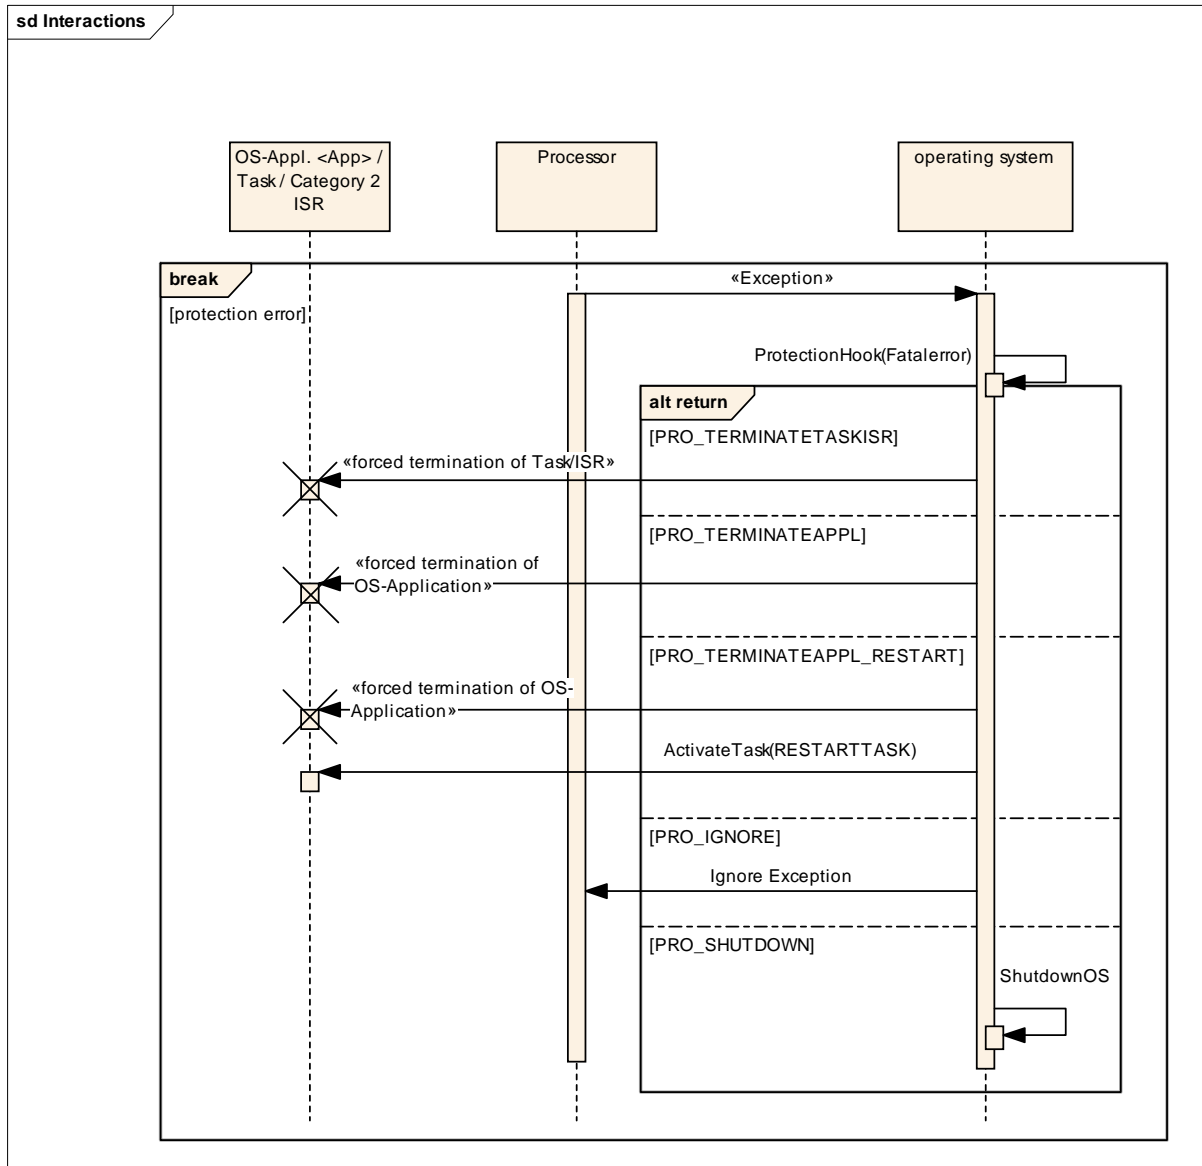


Figure 9.3: Protection Hook sequence chart

The sequence shows the flow of control if a protection error occurs. Depending on the return values of the ProtectionHook, either the faulty Task/ISR is forcibly terminated or the OS-Application is forcibly terminated or the system is shut down. If the action is to terminate the faulty OS-Application an option is to start afterwards the restart task, which can do a cleanup, etc.

9.4 Sequence chart for StartupHook

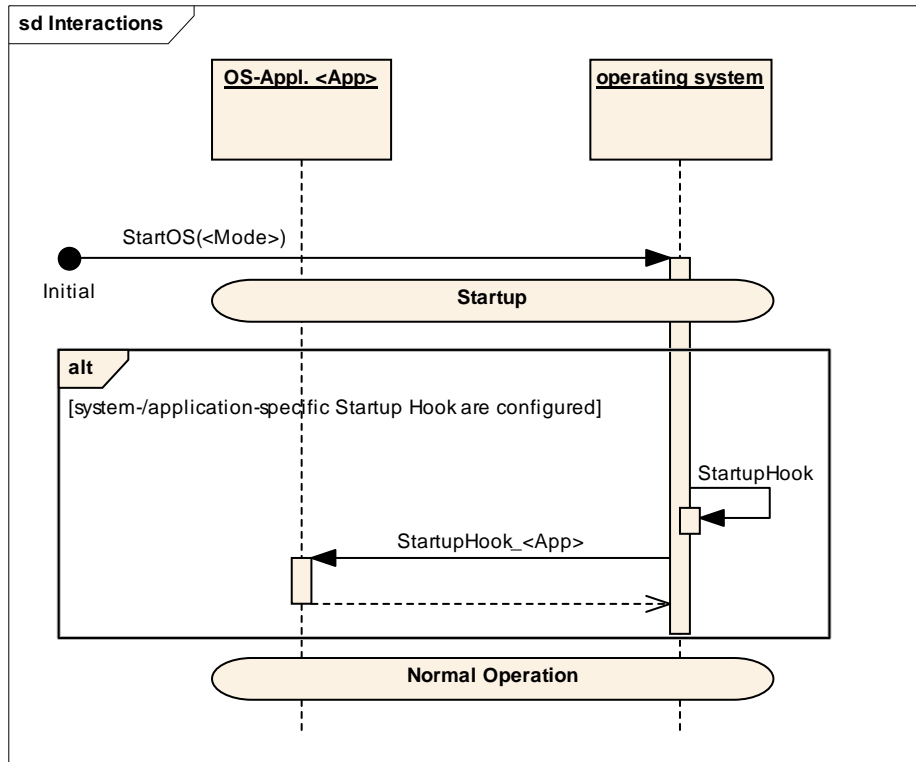


Figure 9.4: StartupHook sequence chart

The above sequence shows the flow of control during the startup of the OS. Like in OSEK OS the user calls the `startOS()` service to start the OS. During the startup the startup hooks are called in the above order. The rest of the startup sequence is identical to the defined behaviour of OSEK OS.

9.5 Sequence chart for ShutdownHook

The next sequence shows the behaviour in case of a shut down. The flow is the same as in OSEK OS with the exception that the shut down hooks of the OS-Applications are called before the general ShutdownHook is called. Note that the specific shutdown hooks of the application are not allowed to block, they must return to the caller.

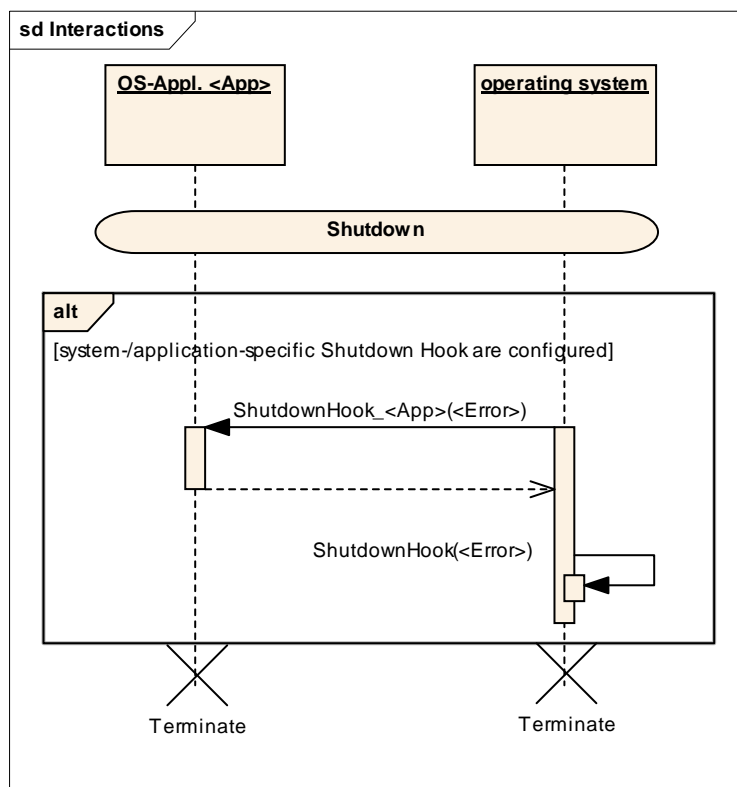


Figure 9.5: ShutdownHook sequence chart

9.6 Sequence diagrams of Sender Receiver communication over the IOC

9.6.1 LastIsBest communication

The figure 11 shows a sequence of successful and failure cases in the interaction between the IOC and the RTE in case of LastIsBest communication (“data” semantic).

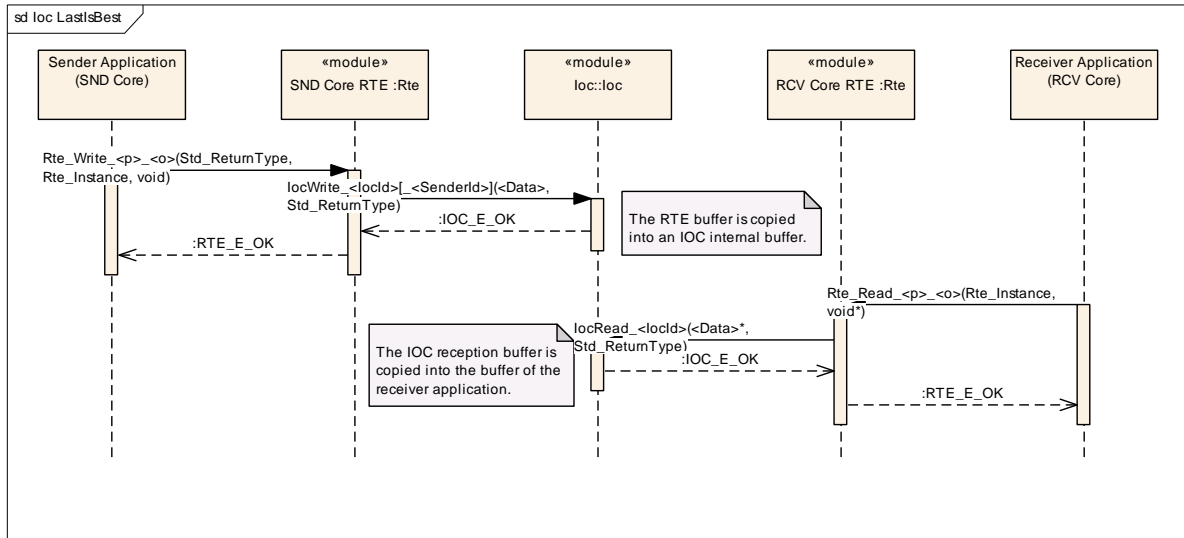


Figure 12: IOC - LastIsBest communication

9.6.2 Queued communication without pull callback

The figure 12 shows the interaction between IOC and RTE with a focus on the congestion control for a queued communication.

The defined communication has no callback functionality for data reception, has an internal buffer size of 2 data elements, no waitpoints are defined and the implicated OS-Applications are located on different cores.

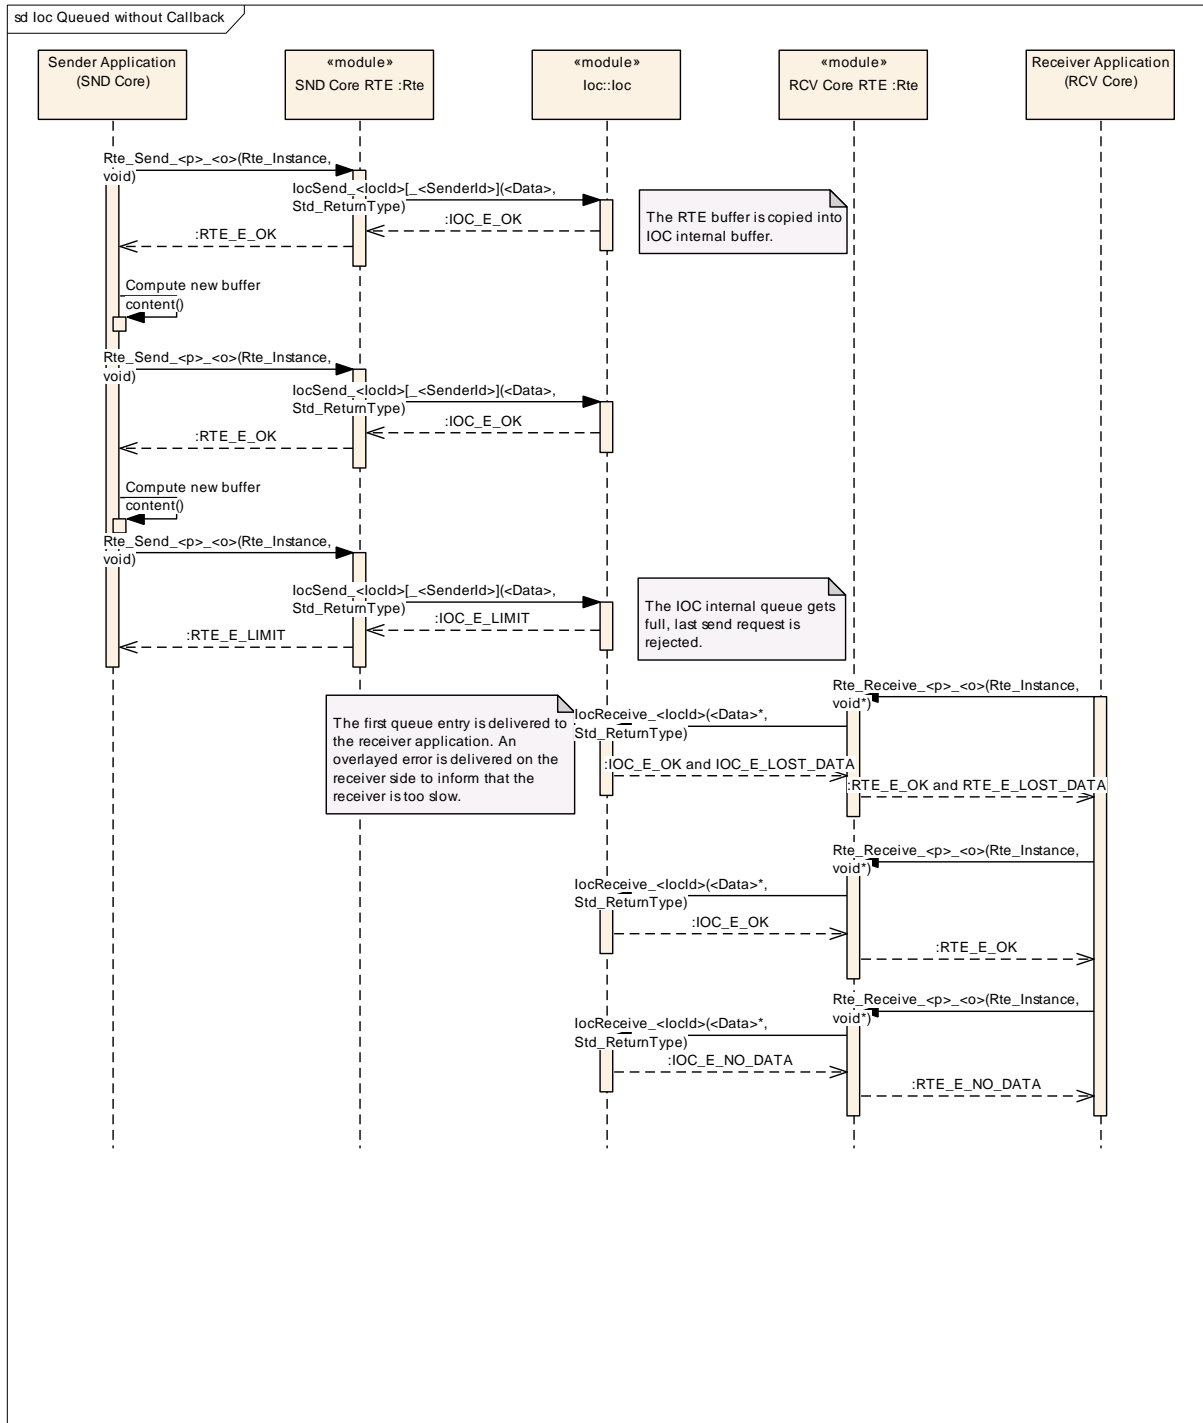


Figure 13: Ioc - Queued communication without callback

9.6.3 Queued communication with pull callback

The figure 13 shows the interaction between IOC and RTE in case of a queued communication with an activated callback functionality. The RTE might handle notification internally and might therefore not provide any callback functions, but a similar scenario will occur in case of communication between CDDs on different cores. The receiving CDD will provide the callback function in this case.

The defined communication has no waitpoints and describes a communication implicating two OS-Applications located on different cores.

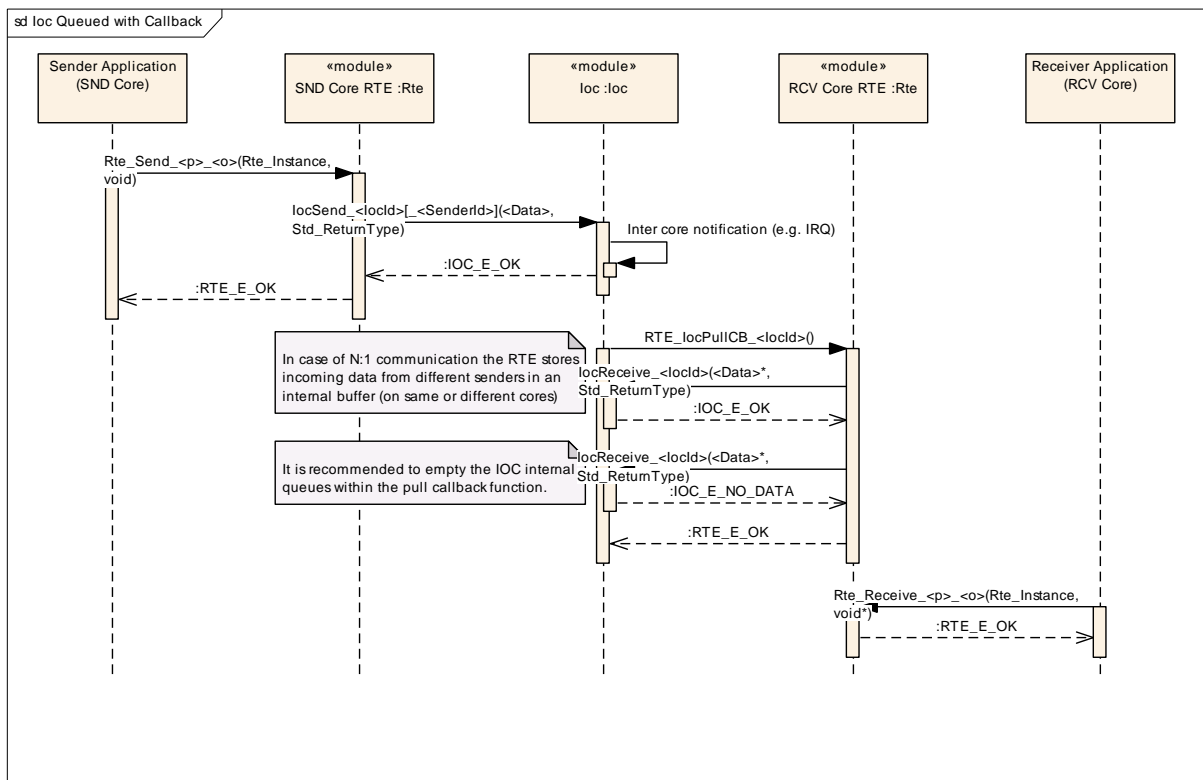


Figure 14: IOC Queued Communication with callback

10 Configuration Specification

In general, this chapter defines configuration parameters and their clustering into containers. In order to support the specification Chapter 10.1 describes fundamentals. It also specifies a template (table) you shall use for the parameter specification.

Chapter 10.2 specifies the structure (containers) and the parameters of the module Os.

Chapter 10.4 specifies published information of the module Os.

10.1 How to read this chapter

In addition to this section, it is highly recommended to read the documents:

- AUTOSAR Layered Software Architecture [1]
- AUTOSAR ECU Configuration Specification [10]
This document describes the AUTOSAR configuration methodology and the AUTOSAR configuration metamodel in detail.

The following is only a short survey of the topic and it will not replace the ECU Configuration Specification document.

10.1.1 Configuration and configuration parameters

Configuration parameters define the variability of the generic part(s) of an implementation of a module. This means that only generic or configurable module implementation can be adapted to the environment (software/hardware) in use during system and/or ECU configuration.

The configuration of parameters can be achieved at different times during the software process: before compile time, before link time or after build time. In the following, the term “configuration class” (of a parameter) shall be used in order to refer to a specific configuration point in time.

Note that not all attributes may be available in all scalability class.

Memory protection configuration is not standardized and therefore not part of this specification.

10.1.2 Variants

Variants describe sets of configuration parameters. E.g., variant 1: only pre-compile time configuration parameters; variant 2: mix of pre-compile- and post build time-configuration parameters. In one variant a parameter can only be of one configuration class.

10.1.3 Containers

Containers structure the set of configuration parameters. This means:

- all configuration parameters are kept in containers.
- (sub-) containers can reference (sub-) containers. It is possible to assign a multiplicity to these references. The multiplicity then defines the possible number of instances of the contained parameters.

10.1.4 Rules for paramters

Some configuration parameters are configured as floating point values and sometimes these values must be rounded in order to be used. The following rules define the rounding of specific parameters:

- Execution times (for the timing protection) are “round down”
- Timeframes are “round down”

10.2 Containers and configuration parameters

The following chapters summarize all configuration parameters and their containers. Background information about the detailed meaning of the parameters can be found in chapters 7 and 8.

For better readability OIL names of the 2.1 OS specification are given in curly braces in the namefield of configuration parameters.

10.2.1 Variants

[OS558] [The configuration of the AUTOSAR OS allows only pre-compile (“VARIANT-PRE-COMPILE“) time configuration parameters.] ()

10.2.2 Os

Module Name	Os
Module Description	Configuration of the Os (Operating System) module.

Included Containers		
Container Name	Multiplicity	Scope / Dependency
OsAlarm	0..*	An OsAlarm may be used to asynchronously inform or activate a specific task. It is possible to start alarms automatically at system start-up depending on the application mode.
OsAppMode	1..*	OsAppMode is the object used to define OSEK OS properties for an OSEK OS application mode. No standard attributes are defined for AppMode. In a CPU, at least one AppMode object has to be defined. [source: OSEK OIL Spec. 2.5] An OsAppMode called OSDEFAULTAPPMODE must always be there for OSEK compatibility.

OsApplication	0..*	An AUTOSAR OS must be capable of supporting a collection of OS objects (tasks, interrupts, alarms, hooks etc.) that form a cohesive functional unit. This collection of objects is termed an OS-Application. All objects which belong to the same OS-Application have access to each other. Access means to allow to use these objects within API services. Access by other applications can be granted separately.
OsCounter	0..*	Configuration information for the counters that belong to the OsApplication.
OsEvent	0..*	Representation of OS events in the configuration context. Adopted from the OSEK OIL specification.
OsIoc	0..1	Configuration of the IOC (Inter OS Application Communicator).
OsIsr	0..*	The OsIsr container represents an OSEK interrupt service routine.
OsOS	1	OS is the object used to define OSEK OS properties for an OSEK application. Per CPU exactly one OS object has to be defined.
OsResource	0..*	An OsResource object is used to coordinate the concurrent access by tasks and ISRs to a shared resource, e.g. the scheduler, any program sequence, memory or any hardware area.
OsScheduleTable	0..*	An OsScheduleTable addresses the synchronization issue by providing an encapsulation of a statically defined set of alarms that cannot be modified at runtime.
OsSpinlock	0..*	An OsSpinlock object is used to coordinate concurrent access by TASKs/ISR2s on different cores to a shared resource.
OsTask	0..*	This container represents an OSEK task.

10.2.3 OsAlarmSetEvent

SWS Item	OS016_Conf :		
Container Name	OsAlarmSetEvent{SETEVENT}		
Description	This container specifies the parameters to set an event		
Configuration Parameters			

SWS Item	OS017_Conf :		
Name	OsAlarmSetEventRef {EVENT}		
Description	Reference to the event that will be set by that alarm action		
Multiplicity	1		
Type	Reference to [OsEvent]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS018_Conf :		
Name	OsAlarmSetEventTaskRef {TASK}		
Description	Reference to the task that will be activated by that event		
Multiplicity	1		
Type	Reference to [OsTask]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

No Included Containers

10.2.4 OsAlarm

SWS Item	OS003_Conf :		
Container Name	OsAlarm{ALARM}		
Description	An OsAlarm may be used to asynchronously inform or activate a specific task. It is possible to start alarms automatically at system start-up depending on the application mode.		
Configuration Parameters			

SWS Item	OS004_Conf :		
Name	OsAlarmAccessingApplication {ACCESSING_APPLICATION}		
Description	Reference to applications which have an access to this object.		
Multiplicity	0..*		
Type	Reference to [OsApplication]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS005_Conf :		
Name	OsAlarmCounterRef {COUNTER}		
Description	Reference to the assigned counter for that alarm		
Multiplicity	1		
Type	Reference to [OsCounter]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

Included Containers		
Container Name	Multiplicity	Scope / Dependency
OsAlarmAction	1	This container defines which type of notification is used when the alarm expires.
OsAlarmAutostart	0..1	If present this container defines if an alarm is started automatically at system start-up depending on the application mode.

10.2.5 OsAlarmAction

SWS Item	OS006_Conf :
Choice container Name	OsAlarmAction{ACTION}
Description	This container defines which type of notification is used when the alarm expires.

Container Choices		
Container Name	Multiplicity	Scope / Dependency
OsAlarmActivateTask	0..1	This container specifies the parameters to activate a task.
OsAlarmCallback	0..1	This container specifies the parameters to call a callback OS alarm action.
OsAlarmIncrementCounter	0..1	This container specifies the parameters to increment a counter.
OsAlarmSetEvent	0..1	This container specifies the parameters to set an event

10.2.6 OsAlarmActivateTask

SWS Item	OS007_Conf :
Container Name	OsAlarmActivateTask{ACTIVATETASK}
Description	This container specifies the parameters to activate a task.
Configuration Parameters	

SWS Item	OS008_Conf :		
Name	OsAlarmActivateTaskRef {TASK}		
Description	Reference to the task that will be activated by that alarm action		
Multiplicity	1		
Type	Reference to [OsTask]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

No Included Containers

10.2.7 OsAlarmAutostart

SWS Item	OS009_Conf :
Container Name	OsAlarmAutostart{AUTOSTART}
Description	If present this container defines if an alarm is started automatically at system start-up depending on the application mode.
Configuration Parameters	

SWS Item	OS010_Conf :		
Name	OsAlarmAlarmTime {ALARMTIME}		
Description	The relative or absolute tick value when the alarm expires for the first time. Note that for an alarm which is RELATIVE the value must be at bigger than 0.		
Multiplicity	1		
Type	EcuIntegerParamDef		
Range	0 .. 18446744073709551615		

Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS011_Conf :		
Name	OsAlarmAutostartType		
Description	This specifies the type of autostart for the alarm..		
Multiplicity	1		
Type	EcucEnumerationParamDef		
Range	ABSOLUTE	The alarm is started on startup via SetAbsAlarm().	
	RELATIVE	The alarm is started on startup via SetRelAlarm().	
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS012_Conf :		
Name	OsAlarmCycleTime {CYCLETIME}		
Description	Cycle time of a cyclic alarm in ticks. If the value is 0 than the alarm is not cyclic.		
Multiplicity	1		
Type	EcucIntegerParamDef		
Range	0 .. 18446744073709551615		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS013_Conf :		
Name	OsAlarmAppModeRef {APPMODE}		
Description	Reference to the application modes for which the AUTOSTART shall be performed		
Multiplicity	1..*		
Type	Reference to [OsAppMode]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

No Included Containers

10.2.8 OsAlarmCallback

SWS Item	OS014_Conf :		
Container Name	OsAlarmCallback{ALARMCALLBACK}		
Description	This container specifies the parameters to call a callback OS alarm action.		
Configuration Parameters			

SWS Item	OS087_Conf :		
Name	OsAlarmCallbackName {ALARMCALLBACKNAME}		
Description	Name of the function that is called when this alarm callback is triggered.		
Multiplicity	1		
Type	EcucFunctionNameDef		
Default value	--		
maxLength	--		
minLength	--		
regularExpression	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

No Included Containers

10.2.9 OsAlarmIncrementCounter

SWS Item	OS302_Conf :		
Container Name	OsAlarmIncrementCounter{INCREMENTCOUNTER}		
Description	This container specifies the parameters to increment a counter.		
Configuration Parameters			

SWS Item	OS015_Conf :		
Name	OsAlarmIncrementCounterRef {COUNTER}		
Description	Reference to the counter that will be incremented by that alarm action		
Multiplicity	1		
Type	Reference to [OsCounter]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU		

No Included Containers

10.2.10 OsApplication

SWS Item	OS114_Conf :		
Container Name	OsApplication{APPLICATION}		
Description	An AUTOSAR OS must be capable of supporting a collection of OS objects (tasks, interrupts, alarms, hooks etc.) that form a cohesive functional unit. This collection of objects is termed an OS-Application. All objects which belong to the same OS-Application have access to each other. Access means to allow to use these objects within API services. Access by other applications can be granted separately.		
Configuration Parameters			

SWS Item	MCOS1020_Conf :		
Name	OsApplicationCoreAssignment {CORE}		
Description	ID of the core onto which the OsApplication is bound.		
Multiplicity	0..1		
Type	EcucIntegerParamDef		
Range	0 .. 65534		

Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: The parameter defines onto which a single OS-Application instance is bound.		

SWS Item	OS115_Conf :		
Name	OsTrusted {TRUSTED}		
Description	Parameter to specify if an OS-Application is trusted or not. true: OS-Application is trusted false: OS-Application is not trusted (default)		
Multiplicity	1		
Type	EcucBooleanParamDef		
Default value	false		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU dependency: Required for scalability class 3 and 4.		

SWS Item	OS231_Conf :		
Name	OsAppAlarmRef		
Description	Specifies the OsAlarms that belong to the OsApplication.		
Multiplicity	0..*		
Type	Reference to [OsAlarm]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU dependency: Required for scalability class 3 and 4		

SWS Item	OS234_Conf :		
Name	OsAppCounterRef		
Description	References the OsCounters that belong to the OsApplication.		
Multiplicity	0..*		
Type	Reference to [OsCounter]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU dependency: Required for scalability class 3 and 4.		

SWS Item	OS392_Conf :		
Name	OsAppEcucPartitionRef		
Description	Denotes which "EcucPartition" is implemented by this "OSApplication".		
Multiplicity	0..1		
Type	Reference to [EcucPartition]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS221_Conf :		
Name	OsApplsRef		
Description	references which Oslrs belong to the OsApplication		

Multiplicity	0..*		
Type	Reference to [OsIsrc]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU dependency: Required for scalability class 3 and 4.		

SWS Item	OS230_Conf :		
Name	OsAppScheduleTableRef		
Description	References the OsScheduleTables that belong to the OsApplication.		
Multiplicity	0..*		
Type	Reference to [OsScheduleTable]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU dependency: Required for scalability class 3 and 4.		

SWS Item	OS116_Conf :		
Name	OsAppTaskRef		
Description	references which OsTasks belong to the OsApplication		
Multiplicity	0..*		
Type	Reference to [OsTask]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU dependency: Required for scalability class 3 and 4		

SWS Item	OS120_Conf :		
Name	OsRestartTask {RESTARTTASK}		
Description	Optionally one task of an OS-Application may be defined as Restart Task. Multiplicity = 1: Restart Task is activated by the Operating System if the protection hook requests it. Multiplicity = 0: No task is automatically started after a protection error happened.		
Multiplicity	0..1		
Type	Reference to [OsTask]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU dependency: Required for scalability class 3 and 4.		

Included Containers		
Container Name	Multiplicity	Scope / Dependency
OsApplicationHooks	1	Container to structure the OS-Application-specific hooks
OsApplicationTrustedFunction	0..*	Container to structure the configuration parameters of trusted functions

10.2.11 OsApplicationHooks

SWS Item	OS020_Conf :		
Container Name	OsApplicationHooks		

Description	Container to structure the OS-Application-specific hooks
Configuration Parameters	

SWS Item	OS213_Conf :		
Name	OsAppErrorHook {ERRORHOOK}		
Description	Select the OS-Application error hook. true: Hook is called false: Hook is not called		
Multiplicity	1		
Type	EcucBooleanParamDef		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU dependency: Required for scalability class 3 and 4.		

SWS Item	OS125_Conf :		
Name	OsAppShutdownHook {SHUTDOWNHOOK}		
Description	Select the OS-Application specific shutdown hook for the OS-Application. true: Hook is called false: Hook is not called		
Multiplicity	1		
Type	EcucBooleanParamDef		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU dependency: Required for scalability class 3 and 4.		

SWS Item	OS124_Conf :		
Name	OsAppStartupHook {STARTUPHOOK}		
Description	Select the OS-Application specific startup hook for the OS-Application. true: Hook is called false: Hook is not called		
Multiplicity	1		
Type	EcucBooleanParamDef		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU dependency: Required for scalability class 3 and 4.		

No Included Containers

10.2.12 OsApplicationTrustedFunction

SWS Item	OS021_Conf :		
Container Name	OsApplicationTrustedFunction		
Description	Container to structure the configuration parameters of trusted functions		
Configuration Parameters			

SWS Item	OS254_Conf :		
Name	OsTrustedFunctionName		
Description	Trusted function (as part of a trusted OS-Application) available to other OS-Applications. This also supersedes the OSEK OIL attribute		

	TRUSTED in APPLICATION because the optionality of this parameter is describing that already.		
Multiplicity	1		
Type	EcucFunctionNameDef		
Default value	--		
maxLength	--		
minLength	--		
regularExpression	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU dependency: Required for scalability class 3 and 4 and in trusted OS-Applications.		

No Included Containers

10.2.13 OsAppMode

SWS Item	OS022_Conf :		
Container Name	OsAppMode{APPMODE}		
Description	<p>OsAppMode is the object used to define OSEK OS properties for an OSEK OS application mode. No standard attributes are defined for AppMode. In a CPU, at least one AppMode object has to be defined. [source: OSEK OIL Spec. 2.5] An OsAppMode called OSDEFAULTAPPMODE must always be there for OSEK compatibility.</p>		
Configuration Parameters			

No Included Containers

10.2.14 OsCounter

SWS Item	OS026_Conf :		
Container Name	OsCounter{COUNTER}		
Description	Configuration information for the counters that belong to the OsApplication.		
Configuration Parameters			

SWS Item	OS027_Conf :		
Name	OsCounterMaxAllowedValue {MAXALLOWEDVALUE}		
Description	Maximum possible allowed value of the system counter in ticks.		
Multiplicity	1		
Type	EcucIntegerParamDef		
Range	1 .. 18446744073709551615		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS028_Conf :		
-----------------	---------------------	--	--

Name	OsCounterMinCycle {MINCYCLE}		
Description	The MINCYCLE attribute specifies the minimum allowed number of counter ticks for a cyclic alarm linked to the counter.		
Multiplicity	1		
Type	EcucIntegerParamDef		
Range	1 .. 18446744073709551615		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS029_Conf :		
Name	OsCounterTicksPerBase {TICKSPERBASE}		
Description	The TICKSPERBASE attribute specifies the number of ticks required to reach a counterspecific unit. The interpretation is implementation-specific.		
Multiplicity	1		
Type	EcucIntegerParamDef		
Range	1 .. 4294967295		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS255_Conf :		
Name	OsCounterType {TYPE}		
Description	This parameter contains the natural type or unit of the counter.		
Multiplicity	1		
Type	EcucEnumerationParamDef		
Range	HARDWARE	This counter is driven by some hardware e.g. a hardware timer unit.	
	SOFTWARE	The counter is driven by some software which calls the IncrementCounter service.	
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU		

SWS Item	OS030_Conf :		
Name	OsSecondsPerTick		
Description	Time of one counter tick in seconds.		
Multiplicity	0..1		
Type	EcucFloatParamDef		
Range	0 .. INF		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU		

SWS Item	OS031_Conf :		
-----------------	---------------------	--	--

Name	OsCounterAccessingApplication {ACCESSING_APPLICATION}		
Description	Reference to applications which have an access to this object.		
Multiplicity	0..*		
Type	Reference to [OsApplication]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

Included Containers		
Container Name	Multiplicity	Scope / Dependency
OsDriver	0..1	This Container contains the information who will drive the counter. This configuration is only valid if the counter has OsCounterType set to HARDWARE. If the container does not exist (multiplicity=0) the timer is managed by the OS internally (OSINTERNAL). If the container exists the OS can use the GPT interface to manage the timer. The user have to supply the GPT channel. If the counter is driven by some other (external to the OS) source (like a TPU for example) this must be described as a vendor specific extension.
OsTimeConstant	0..*	Allows the user to define constants which can be e.g. used to compare time values with timer tick values. A time value will be converted to a timer tick value during generation and can later on accessed via the OsConstName. The conversation is done by rounding time values to the nearest fitting tick value.

10.2.15 OsEvent

SWS Item	OS033_Conf :
Container Name	OsEvent{EVENT}
Description	Representation of OS events in the configuration context. Adopted from the OSEK OIL specification.
Configuration Parameters	

SWS Item	OS034_Conf :		
Name	OsEventMask {MASK}		
Description	If event mask would be set to AUTO in OIL, this parameter should be omitted here.		
Multiplicity	0..1		
Type	EcucIntegerParamDef		
Range	0 .. 18446744073709551615		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

No Included Containers

10.2.16 OsHooks

SWS Item	OS035_Conf :		
Container Name	OsHooks		
Description	Container to structure all hooks belonging to the OS		
Configuration Parameters			

SWS Item	OS036_Conf :		
Name	OsErrorHook {ERRORHOOK}		
Description	Error hook as defined by OSEK true: Hook is called false: Hook is not called		
Multiplicity	1		
Type	EcucBooleanParamDef		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS037_Conf :		
Name	OsPostTaskHook {POSTTASKHOOK}		
Description	Post-task hook as defined by OSEK true: Hook is called false: Hook is not called		
Multiplicity	1		
Type	EcucBooleanParamDef		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS038_Conf :		
Name	OsPreTaskHook {PRETASKHOOK}		
Description	Pre-task hook as defined by OSEK true: Hook is called false: Hook is not called		
Multiplicity	1		
Type	EcucBooleanParamDef		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS214_Conf :		
Name	OsProtectionHook {PROTECTIONHOOK}		
Description	Switch to enable/disable the call to the (user supplied) protection hook. true: Protection hook is called on protection error false: Protection hook is not called		
Multiplicity	0..1		
Type	EcucBooleanParamDef		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	

	Post-build time	--	
Scope / Dependency	scope: ECU dependency: Required for scalability class 2,3 and 4		

SWS Item	OS039_Conf :		
Name	OsShutdownHook {SHUTDOWNHOOK}		
Description	Shutdown hook as defined by OSEK true: Hook is called false: Hook is not called		
Multiplicity	1		
Type	EcucBooleanParamDef		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS040_Conf :		
Name	OsStartupHook {STARTUPHOOK}		
Description	Startup hook as defined by OSEK true: Hook is called false: Hook is not called		
Multiplicity	1		
Type	EcucBooleanParamDef		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

No Included Containers

10.2.17 Oslsr

SWS Item	OS041_Conf :		
Container Name	Oslsr{ISR}		
Description	The Oslsr container represents an OSEK interrupt service routine.		
Configuration Parameters			

SWS Item	OS042_Conf :		
Name	OslsrCategory {CATEGORY}		
Description	This attribute specifies the category of this ISR.		
Multiplicity	1		
Type	EcucEnumerationParamDef		
Range	CATEGORY_1	Interrupt is of category 1	
	CATEGORY_2	Interrupt is of category 2	
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS043_Conf :		
Name	OslsrResourceRef {RESOURCE}		
Description	This reference defines the resources accessed by this ISR.		
Multiplicity	0..*		

Type	Reference to [OsResource]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

Included Containers		
Container Name	Multiplicity	Scope / Dependency
OsIsrTimingProtection	0..1	This container contains all parameters which are related to timing protection. If the container exists, the timing protection is used for this interrupt. If the container does not exist, the interrupt is not supervised regarding timing violations.

10.2.18 OsIsrResourceLock

SWS Item	OS388_Conf :
Container Name	OsIsrResourceLock{LOCKINGTIME}
Description	This container contains a list of times the interrupt uses resources.
Configuration Parameters	

SWS Item	OS389_Conf :		
Name	OsIsrResourceLockBudget {MAXRESOURCELOCKTIME}		
Description	This parameter contains the maximum time the interrupt is allowed to hold the given resource (in seconds).		
Multiplicity	1		
Type	EcucFloatParamDef		
Range	0 .. INF		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU dependency: Required for scalability class 2 and 4		

SWS Item	OS390_Conf :		
Name	OsIsrResourceLockResourceRef {RESOURCE}		
Description	Reference to the resource the locking time is depending on		
Multiplicity	1		
Type	Reference to [OsResource]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU dependency: Required for scalability class 2 and 4		

No Included Containers

10.2.19 OsIsrTimingProtection

SWS Item	OS326_Conf :
Container Name	OsIsrTimingProtection{TIMING_PROTECTION}

Description	This container contains all parameters which are related to timing protection. If the container exists, the timing protection is used for this interrupt. If the container does not exist, the interrupt is not supervised regarding timing violations.
Configuration Parameters	

SWS Item	OS229_Conf :		
Name	OslsrAllInterruptLockBudget {MAXALLINTERRUPTLOCKTIME}		
Description	This parameter contains the maximum time for which the ISR is allowed to lock all interrupts (via SuspendAllInterrupts() or DisableAllInterrupts()) (in seconds).		
Multiplicity	0..1		
Type	EcucFloatParamDef		
Range	0 .. INF		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU dependency: Required for scalability class 2 and 4		

SWS Item	OS222_Conf :		
Name	OslsrExecutionBudget {EXECUTIONTIME}		
Description	The parameter contains the maximum allowed execution time of the interrupt (in seconds).		
Multiplicity	0..1		
Type	EcucFloatParamDef		
Range	0 .. INF		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU dependency: Required for scalability class 2 and 4		

SWS Item	OS387_Conf :		
Name	OslsrOSInterruptLockBudget {MAXOSINTERRUPTLOCKTIME}		
Description	This parameter contains the maximum time for which the ISR is allowed to lock all Category 2 interrupts (via SuspendOSInterrupts()) (in seconds).		
Multiplicity	0..1		
Type	EcucFloatParamDef		
Range	0 .. INF		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU dependency: Required for scalability class 2 and 4		

SWS Item	OS223_Conf :		
Name	OslsrTimeFrame {TIMEFRAME}		
Description	This parameter contains the minimum inter-arrival time between successive interrupts (in seconds).		
Multiplicity	0..1		

Type	EcucFloatParamDef		
Range	0 .. INF		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU dependency: Required for scalability class 2 and 4		

Included Containers		
Container Name	Multiplicity	Scope / Dependency
OsIsrResourceLock	0..*	This container contains a list of times the interrupt uses resources.

10.2.20 OsOS

SWS Item	OS044_Conf :
Container Name	OsOS{OS}
Description	OS is the object used to define OSEK OS properties for an OSEK application. Per CPU exactly one OS object has to be defined.
Configuration Parameters	

SWS Item	MCOS1019_Conf :		
Name	OsNumberOfCores		
Description	Maximum number of cores that are controlled by the OS. The OS uses the value internally. It depends on the ECU HW.		
Multiplicity	0..1		
Type	EcucIntegerParamDef		
Range	1 .. 65535		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS259_Conf :		
Name	OsScalabilityClass {SCALABILITYCLASS}		
Description	A scalability class for each System Object "OS" has to be selected. In order to customize the operating system to the needs of the user and to take full advantage of the processor features the operating system can be scaled according to the scalability classes. If the scalability class is omitted this translates to the OIL AUTO mechanism.		
Multiplicity	0..1		
Type	EcucEnumerationParamDef		
Range	SC1	--	
	SC2	--	
	SC3	--	
	SC4	--	
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU		

SWS Item	OS307_Conf :
-----------------	---------------------

Name	OsStackMonitoring {STACKMONITORING}		
Description	Select stack monitoring of Tasks/Category 2 ISRs true: Stacks are monitored false: Stacks are not monitored		
Multiplicity	1		
Type	EcucBooleanParamDef		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU		

SWS Item	OS046_Conf :		
Name	OsStatus {STATUS}		
Description	The Status attribute specifies whether a system with standard or extended status has to be used. Automatic assignment is not supported for this attribute.		
Multiplicity	1		
Type	EcucEnumerationParamDef		
Range	EXTENDED	--	
	STANDARD	--	
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS047_Conf :		
Name	OsUseGetServiceId {USEGETSERVICEID}		
Description	As defined by OSEK		
Multiplicity	1		
Type	EcucBooleanParamDef		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS048_Conf :		
Name	OsUseParameterAccess {USEPARAMETERACCESS}		
Description	As defined by OSEK		
Multiplicity	1		
Type	EcucBooleanParamDef		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS049_Conf :		
Name	OsUseResScheduler {USERESSCHEDULER}		
Description	The OsUseResScheduler attribute defines whether the resource RES_SCHEDULER is used within the application.		
Multiplicity	1		
Type	EcucBooleanParamDef		
Default value	true		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	

	Link time	--	
	Post-build time	--	
Scope / Dependency			

Included Containers		
Container Name	Multiplicity	Scope / Dependency
OsHooks	1	Container to structure all hooks belonging to the OS

10.2.21 OsResource

SWS Item	OS252_Conf :		
Container Name	OsResource{RESOURCE}		
Description	An OsResource object is used to co-ordinate the concurrent access by tasks and ISRs to a shared resource, e.g. the scheduler, any program sequence, memory or any hardware area.		
Configuration Parameters			

SWS Item	OS050_Conf :		
Name	OsResourceProperty {RESOURCEPROPERTY}		
Description	This specifies the type of the resource.		
Multiplicity	1		
Type	EcucEnumerationParamDef		
Range	INTERNAL	The resource is an internal resource.	
	LINKED	The resource is a linked resource (a second name for a existing resource).	
	STANDARD	The resource is a standard resource.	
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS051_Conf :		
Name	OsResourceAccessingApplication {ACCESSING_APPLICATION}		
Description	Reference to applications which have an access to this object.		
Multiplicity	0..*		
Type	Reference to [OsApplication]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS052_Conf :		
Name	OsResourceLinkedResourceRef {LINKEDRESOURCE}		
Description	The link to the resource. Must be valid if OsResourceProperty is LINKED. If OsResourceProperty is not LINKED the value is ignored.		
Multiplicity	0..1		
Type	Reference to [OsResource]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

No Included Containers

10.2.22 OsScheduleTable

SWS Item	OS141_Conf :
Container Name	OsScheduleTable{SCHEDULETABLE}
Description	An OsScheduleTable addresses the synchronization issue by providing an encapsulation of a statically defined set of alarms that cannot be modified at runtime.
Configuration Parameters	

SWS Item	OS053_Conf :		
Name	OsScheduleTableDuration		
Description	This parameter defines the modulus of the schedule table (in ticks).		
Multiplicity	1		
Type	EcucIntegerParamDef		
Range	0 .. 18446744073709551615		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS144_Conf :		
Name	OsScheduleTableRepeating {REPEATING}		
Description	true: first expiry point on the schedule table shall be processed at final expiry point delay ticks after the final expiry point is processed. false: the schedule table processing stops when the final expiry point is processed.		
Multiplicity	1		
Type	EcucBooleanParamDef		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU		

SWS Item	OS054_Conf :		
Name	OsSchTblAccessingApplication {ACCESSING_APPLICATION}		
Description	Reference to applications which have an access to this object.		
Multiplicity	0..*		
Type	Reference to [OsApplication]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS145_Conf :
Name	OsScheduleTableCounterRef {COUNTER}
Description	This parameter contains a reference to the counter which drives the schedule table.
Multiplicity	1
Type	Reference to [OsCounter]

ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU		

Included Containers		
Container Name	Multiplicity	Scope / Dependency
OsScheduleTableAutostart	0..1	This container specifies if and how the schedule table is started on startup of the Operating System. The options to start a schedule table correspond to the API calls to start schedule tables during runtime.
OsScheduleTableExpiryPoint	1..*	The point on a Schedule Table at which the OS activates tasks and/or sets events
OsScheduleTableSync	0..1	This container specifies the synchronization parameters of the schedule table.

10.2.23 OsScheduleTableAutostart

SWS Item	OS335_Conf :
Container Name	OsScheduleTableAutostart{AUTOSTART}
Description	This container specifies if and how the schedule table is started on startup of the Operating System. The options to start a schedule table correspond to the API calls to start schedule tables during runtime.
Configuration Parameters	

SWS Item	OS056_Conf :	
Name	OsScheduleTableAutostartType	
Description	This specifies the type of the autostart for the schedule table.	
Multiplicity	1	
Type	EcucEnumerationParamDef	
Range	ABSOLUTE	The schedule table is started during startup with the StartScheduleTableAbs() service.
	RELATIVE	The schedule table is started during startup with the StartScheduleTableRel() service.
	SYNCHRON	The schedule table is started during startup with the StartScheduleTableSynchron() service.
ConfigurationClass	Pre-compile time	X All Variants
	Link time	--
	Post-build time	--
Scope / Dependency		

SWS Item	OS057_Conf :	
Name	OsScheduleTableStartValue	
Description	Absolute autostart tick value when the schedule table starts. Only used if the OsScheduleTableAutostartType is ABSOLUTE. Relative offset in ticks when the schedule table starts. Only used if the OsScheduleTableAutostartType is RELATIVE.	
Multiplicity	0..1	
Type	EcucIntegerParamDef	
Range	0 .. 18446744073709551615	
Default value	--	
ConfigurationClass	Pre-compile time	X All Variants

	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU		

SWS Item	OS058_Conf :		
Name	OsScheduleTableAppModeRef		
Description	Reference in which application modes the schedule table should be started during startup		
Multiplicity	1..*		
Type	Reference to [OsAppMode]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU		

No Included Containers

10.2.24 OsScheduleTableEventSetting

SWS Item	OS059_Conf :		
Container Name	OsScheduleTableEventSetting{SETEVENT}		
Description	Event that is triggered by that schedule table.		
Configuration Parameters			

SWS Item	OS060_Conf :		
Name	OsScheduleTableSetEventRef {EVENT}		
Description	Reference to event that will be set by action		
Multiplicity	1		
Type	Reference to [OsEvent]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS061_Conf :		
Name	OsScheduleTableSetEventTaskRef		
Description	--		
Multiplicity	1		
Type	Reference to [OsTask]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

No Included Containers

10.2.25 OsScheduleTableExpiryPoint

SWS Item	OS143_Conf :		
Container Name	OsScheduleTableExpiryPoint{ACTION}		
Description	The point on a Schedule Table at which the OS activates tasks and/or sets events		

Configuration Parameters

SWS Item	OS062_Conf :		
Name	OsScheduleTblExpPointOffset		
Description	The offset from zero (in ticks) at which the expiry point is to be processed.		
Multiplicity	1		
Type	EcuIntegerParamDef		
Range	0 .. 18446744073709551615		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

Included Containers

Container Name	Multiplicity	Scope / Dependency
OsScheduleTableEventSetting	0..*	Event that is triggered by that schedule table.
OsScheduleTableTaskActivation	0..*	Task that is triggered by that schedule table.
OsScheduleTblAdjustableExpPoint	0..1	Adjustable expiry point

10.2.26 OsScheduleTableTaskActivation

SWS Item	OS066_Conf :		
Container Name	OsScheduleTableTaskActivation{ACTIVATETASK}		
Description	Task that is triggered by that schedule table.		
Configuration Parameters			

SWS Item	OS067_Conf :		
Name	OsScheduleTableActivateTaskRef {TASK}		
Description	Reference to task that will be activated by action		
Multiplicity	1		
Type	Reference to [OsTask]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU		

No Included Containers

10.2.27 OsScheduleTblAdjustableExpPoint

SWS Item	OS068_Conf :		
Container Name	OsScheduleTblAdjustableExpPoint		
Description	Adjustable expiry point		
Configuration Parameters			

SWS Item	OS069_Conf :		
Name	OsScheduleTableMaxLengthen		
Description	The maximum positive adjustment that can be made to the expiry point offset (in ticks).		

Multiplicity	1		
Type	EcucIntegerParamDef		
Range	0 .. 18446744073709551615		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS070_Conf :		
Name	OsScheduleTableMaxShorten		
Description	The maximum negative adjustment that can be made to the expiry point offset (in ticks).		
Multiplicity	1		
Type	EcucIntegerParamDef		
Range	0 .. 18446744073709551615		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

No Included Containers

10.2.28 OsScheduleTableSync

SWS Item	OS063_Conf :		
Container Name	OsScheduleTableSync{LOCAL_TO_GLOBAL_TIME_SYNCHRONIZATION}		
Description	This container specifies the synchronization parameters of the schedule table.		
Configuration Parameters			

SWS Item	OS064_Conf :		
Name	OsScheduleTblExplicitPrecision		
Description	This configuration is only valid if the explicit synchronisation is used.		
Multiplicity	0..1		
Type	EcucIntegerParamDef		
Range	0 .. 18446744073709551615		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: System		

SWS Item	OS065_Conf :		
Name	OsScheduleTblSyncStrategy		
Description	AUTOSAR OS provides support for synchronisation in two ways: explicit and implicit.		
Multiplicity	1		
Type	EcucEnumerationParamDef		
Range	EXPLICIT	The schedule table is driven by an OS counter but processing needs to be synchronized with a different counter which is not an OS counter object.	

	IMPLICIT	The counter driving the schedule table is the counter with which synchronisation is required.	
	NONE	No support for synchronisation. (default)	
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: System		

No Included Containers

10.2.29 OsSpinlock

SWS Item	MCOS0258_Conf :		
Container Name	OsSpinlock{SPINLOCK}		
Description	An OsSpinlock object is used to co-ordinate concurrent access by TASKs/ISR2s on different cores to a shared resource.		
Configuration Parameters			

SWS Item	MCOS1021_Conf :		
Name	OsSpinlockAccessingApplication {ACCESSING_APPLICATION}		
Description	Reference to OsApplications that have an access to this object.		
Multiplicity	1..*		
Type	Reference to [OsApplication]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	MCOS1022_Conf :		
Name	OsSpinlockSuccessor {NEXT_SPINLOCK}		
Description	Reference to OsApplications that have an access to this object. To check whether a spinlock can be occupied (in a nested way) without any danger of deadlock, a linked list of spinlocks can be defined. A spinlock can only be occupied in the order of the linked list. It is allowed to skip a spinlock. If no linked list is specified, spinlocks cannot be nested.		
Multiplicity	0..1		
Type	Reference to [OsSpinlock]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

No Included Containers

10.2.30 OsTask

SWS Item	OS073_Conf :
Container Name	OsTask{TASK}
Description	This container represents an OSEK task.
Configuration Parameters	

SWS Item	OS074_Conf :		
Name	OsTaskActivation {ACTIVATION}		
Description	This attribute defines the maximum number of queued activation requests for the task. A value equal to "1" means that at any time only a single activation is permitted for this task. Note that the value must be a natural number starting at 1.		
Multiplicity	1		
Type	EcucIntegerParamDef		
Range	1 .. 4294967295		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS075_Conf :		
Name	OsTaskPriority {PRIORITY}		
Description	The priority of a task is defined by the value of this attribute. This value has to be understood as a relative value, i.e. the values show only the relative ordering of the tasks. OSEK OS defines the lowest priority as zero (0); larger values correspond to higher priorities.		
Multiplicity	1		
Type	EcucIntegerParamDef		
Range	0 .. 4294967295		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS076_Conf :		
Name	OsTaskSchedule {SCHEDULE}		
Description	The OsTaskSchedule attribute defines the preemptability of the task. If this attribute is set to NON, no internal resources may be assigned to this task.		
Multiplicity	1		
Type	EcucEnumerationParamDef		
Range	FULL	Task is preemptable.	
	NON	Task is not preemptable.	
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS077_Conf :		
Name	OsTaskAccessingApplication {ACCESSING_APPLICATION}		
Description	Reference to applications which have an access to this object.		
Multiplicity	0..*		

Type	Reference to [OsApplication]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS078_Conf :		
Name	OsTaskEventRef {EVENT}		
Description	This reference defines the list of events the extended task may react on.		
Multiplicity	0..*		
Type	Reference to [OsEvent]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	OS079_Conf :		
Name	OsTaskResourceRef {RESOURCE}		
Description	This reference defines a list of resources accessed by this task.		
Multiplicity	0..*		
Type	Reference to [OsResource]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

Included Containers		
Container Name	Multiplicity	Scope / Dependency
OsTaskAutostart	0..1	This container determines whether the task is activated during the system start-up procedure or not for some specific application modes. If the task shall be activated during the system start-up, this container is present and holds the references to the application modes in which the task is auto-started.
OsTaskTimingProtection	0..1	This container contains all parameters regarding timing protection of the task.

10.2.31 OsTaskAutostart

SWS Item	OS080_Conf :		
Container Name	OsTaskAutostart{AUTOSTART}		
Description	This container determines whether the task is activated during the system start-up procedure or not for some specific application modes. If the task shall be activated during the system start-up, this container is present and holds the references to the application modes in which the task is auto-started.		
Configuration Parameters			

SWS Item	OS081_Conf :		
Name	OsTaskAppModeRef {APPMODE}		
Description	Reference to application modes in which that task is activated on		

	startup of the OS		
Multiplicity	1..*		
Type	Reference to [OsAppMode]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

No Included Containers

10.2.32 OsTaskResourceLock

SWS Item	OS082_Conf :		
Container Name	OsTaskResourceLock{RESOURCELOCK}		
Description	This container contains the worst case time between getting and releasing a given resource (in seconds).		
Configuration Parameters			

SWS Item	OS083_Conf :		
Name	OsTaskResourceLockBudget {RESOURCELOCKTIME}		
Description	This parameter contains the maximum time the task is allowed to lock the resource (in seconds)		
Multiplicity	1		
Type	EcucFloatParamDef		
Range	0 .. INF		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU dependency: Required for scalability class 2 and 4		

SWS Item	OS084_Conf :		
Name	OsTaskResourceLockResourceRef {RESOURCE}		
Description	Reference to the resource used by the task		
Multiplicity	1		
Type	Reference to [OsResource]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU dependency: Required for scalability class 2 and 4		

No Included Containers

10.2.33 OsTaskTimingProtection

SWS Item	OS325_Conf :		
Container Name	OsTaskTimingProtection{TIMING_PROTECTION}		
Description	This container contains all parameters regarding timing protection of the task.		
Configuration Parameters			

SWS Item	OS085_Conf :		
-----------------	---------------------	--	--

Name	OsTaskAllInterruptLockBudget {MAXALLINTERRUPTLOCKTIME}		
Description	This parameter contains the maximum time for which the task is allowed to lock all interrupts (via SuspendAllInterrupts() or DisableAllInterrupts()) (in seconds).		
Multiplicity	0..1		
Type	EcucFloatParamDef		
Range	0 .. INF		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU dependency: Required for scalability class 2 and 4		

SWS Item	OS185_Conf :		
Name	OsTaskExecutionBudget {EXECUTIONBUDGET}		
Description	This parameter contains the maximum allowed execution time of the task (in seconds).		
Multiplicity	0..1		
Type	EcucFloatParamDef		
Range	0 .. INF		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU dependency: Required for scalability class 2 and 4		

SWS Item	OS086_Conf :		
Name	OsTaskOsInterruptLockBudget {MAXOSINTERRUPTLOCKTIME}		
Description	This parameter contains the maximum time for which the task is allowed to lock all Category 2 interrupts (via SuspendOSInterrupts()) (in seconds).		
Multiplicity	0..1		
Type	EcucFloatParamDef		
Range	0 .. INF		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU dependency: Required for scalability class 2 and 4		

SWS Item	OS391_Conf :		
Name	OsTaskTimeFrame {TIMEFRAME}		
Description	The minimum inter-arrival time between activations and/or releases of a task (in seconds).		
Multiplicity	0..1		
Type	EcucFloatParamDef		
Range	0 .. INF		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency	scope: ECU		

dependency: Only available in scalability class 2 and 4

Included Containers

Container Name	Multiplicity	Scope / Dependency
OsTaskResourceLock	0..*	This container contains the worst case time between getting and releasing a given resource (in seconds).

10.2.34 OsTimeConstant

SWS Item	OS386_Conf :
Container Name	OsTimeConstant{TIMECONSTANTS}
Description	<p>Allows the user to define constants which can be e.g. used to compare time values with timer tick values.</p> <p>A time value will be converted to a timer tick value during generation and can later on accessed via the OsConstName. The conversation is done by rounding time values to the nearest fitting tick value.</p>
Configuration Parameters	

SWS Item	OS002_Conf :									
Name	OsTimeValue									
Description	This parameter contains the value of the constant in seconds.									
Multiplicity	1									
Type	EcucFloatParamDef									
Range	0 .. INF									
Default value	--									
ConfigurationClass	<table border="1"> <tr> <td>Pre-compile time</td> <td>X</td> <td>All Variants</td> </tr> <tr> <td>Link time</td> <td>--</td> <td></td> </tr> <tr> <td>Post-build time</td> <td>--</td> <td></td> </tr> </table>	Pre-compile time	X	All Variants	Link time	--		Post-build time	--	
Pre-compile time	X	All Variants								
Link time	--									
Post-build time	--									
Scope / Dependency	scope: ECU									

No Included Containers

10.3 Containers and configuration parameter extensions of the IOC

This section describes the content of the IOC Configuration Description that is needed for the generation of the IOC API.

10.3.1 Osloc

SWS Item	MCOS1000_Conf :
Container Name	Osloc
Description	Configuration of the IOC (Inter OS Application Communicator).
Configuration Parameters	

Included Containers

Container Name	Multiplicity	Scope / Dependency
OslocCommunication	0..*	Representation of a 1:1 or N:1

	<p>communication between software parts located in different OS-Applications that are bound to the same or to different cores. The name shall begin with the name of the sending software service and be followed by a unique identifier delivered by the sending software service. In the case of RTE as user attention shall be paid on the fact that uniqueness for identifier names has to be reached over ports, data elements, object instances and maybe additional identification properties (E.g. Case 1:N mapping to 1:1). Example: - <NameSpace>_UniqueID</p>
--	--

10.3.2 OslocCommunication

SWS Item	MCOS1003_Conf :		
Container Name	OslocCommunication		
Description	<p>Representation of a 1:1 or N:1 communication between software parts located in different OS-Applications that are bound to the same or to different cores. The name shall begin with the name of the sending software service and be followed by a unique identifier delivered by the sending software service. In the case of RTE as user attention shall be paid on the fact that uniqueness for identifier names has to be reached over ports, data elements, object instances and maybe additional identification properties (E.g. Case 1:N mapping to 1:1). Example: - <NameSpace>_UniqueID</p>		
Configuration Parameters			

SWS Item	MCOS1001_Conf :		
Name	OslocBufferLength		
Description	<p>This attribute defines the size of the IOC internal queue to be allocated for a queued communication. This configuration information shall allow the optimization of the needed memory for communications requiring buffers within the RTE and within the IOC.</p>		
Multiplicity	0..1		
Type	EcucIntegerParamDef		
Range	0 .. 4294967295		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

Included Containers		
Container Name	Multiplicity	Scope / Dependency
OslocDataProperties	1..*	Data properties of the data to be transferred on the IOC communication channel.
OslocReceiverProperties	1	Representation of receiver properties for

		one communication. For each OslocCommunication (1:1 or N:1) one receiver has to be defined. This container should be instantiated within an OslocCommunication.
OslocSenderProperties	1..*	Representation of sender properties for one communication. For each OslocCommunication one (1:1) or many senders (N:1) have to be defined. Multiplicity > 1 (N:1 communication) is only allowed for Multiplicity of OslocDataTypeRef = 1. This container should be instantiated within an OslocCommunication.

10.3.3 OslocSenderProperties

SWS Item	MCOS1015_Conf :	
Container Name	OslocSenderProperties	
Description	Representation of sender properties for one communication. For each OslocCommunication one (1:1) or many senders (N:1) have to be defined. Multiplicity > 1 (N:1 communication) is only allowed for Multiplicity of OslocDataTypeRef = 1. This container should be instantiated within an OslocCommunication.	
Configuration Parameters		

SWS Item	MCOS1036_Conf :	
Name	OslocFunctionImplementationKind	
Description	This parameter is used to select whether this communication is implemented as a macro or as a function.	
Multiplicity	0..1	
Type	EcucEnumerationParamDef	
Range	DO_NOT_CARE	It is not defined whether a macro or a function is used. (default)
	FUNCTION	Communication is implemented as a function
	MACRO	Communication is implemented as a macro
ConfigurationClass	Pre-compile time	X All Variants
	Link time	--
	Post-build time	--
Scope / Dependency		

SWS Item	MCOS1016_Conf :	
Name	OslocSenderId	
Description	Representation of a sender in a N:1 communication to distinguish between senders. This parameter does not exist in 1:1 communication.	
Multiplicity	0..1	
Type	EcucIntegerParamDef	
Range	0 .. 255	
Default value	--	
ConfigurationClass	Pre-compile time	X All Variants
	Link time	--
	Post-build time	--

Scope / Dependency	scope: This identifier shall be used for the generation of the function name for the sending of data (See API definition for functions)
---------------------------	---

SWS Item	MCOS1014_Conf :		
Name	OslocSendingOsApplicationRef		
Description	This attribute is a reference to the sending OS-Application instance defined in the configuration file of the OS. This information shall allows the generator to get additional information necessary for the code generation like: * The protection properties of the communicating OS-Applications to find out which protection boundaries have to be crossed. * The core identifiers to find out if an intra or an inter core communication has to be realized * Interrupt details in case of cross core notification to realize over IRQs		
Multiplicity	1		
Type	Reference to [OsApplication]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

No Included Containers

10.3.4 OslocReceiverProperties

SWS Item	MCOS1017_Conf :		
Container Name	OslocReceiverProperties		
Description	Representation of receiver properties for one communication. For each OslocCommunication (1:1 or N:1) one receiver has to be defined. This container should be instanciated within an OslocCommunication.		
Configuration Parameters			

SWS Item	MCOS1036_Conf :		
Name	OslocFunctionImplementationKind		
Description	This parameter is used to select whether this communication is implemented as a macro or as a function.		
Multiplicity	0..1		
Type	EcucEnumerationParamDef		
Range	DO_NOT_CARE		It is not defined whether a macro or a function is used. (default)
	FUNCTION		Communication is implemented as a function
	MACRO		Communication is implemented as a macro
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	MCOS1010_Conf :		
Name	OslocReceiverPullCB		
Description	This attribute defines the name of a callback function that the IOC shall call on the receiving core for each data reception. In case of non existence of this attribute no ReceiverPullCB notification shall be		

	applied by the IOC. The name of the function shall begin with the name of the receiving module, followed with a callback name and followed by the locId. Example: void RTE_ReceiverPullCB_RTE25 (void). If this attribute does not exist, it means that no ReceiverPullCB shall be called (No notification from IOC is required). If this attribute exists the IOC shall call the callback function on the receiving core.		
Multiplicity	0..1		
Type	EcucFunctionNameDef		
Default value	--		
maxLength	--		
minLength	--		
regularExpression	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	MCOS1012_Conf :		
Name	OslocReceivingOsApplicationRef		
Description	This attribute is a reference to the receiving OsApplication instance defined in the configuration file of the OS. This information allows for the generator to get additional information necessary for the code generation like: * The protection properties of the communicating OsApplications to find out which protections have to be crossed * The core identifiers to find out if an intra or an inter core communication has to be realized * Interrupt details in case of cross core notification to realize over IRQs		
Multiplicity	1		
Type	Reference to [OsApplication]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

No Included Containers

10.3.5 OslocDataProperties

SWS Item	MCOS1023_Conf :		
Container Name	OslocDataProperties		
Description	Data properties of the data to be transferred on the IOC communication channel.		
Configuration Parameters			

SWS Item	MCOS1035_Conf :		
Name	OslocDataPropertyIndex		
Description	This parameter is used to define in which order the data is send, e.g. whether locSendGroup(A,B) or locSendGroup(B,A) shall be used.		
Multiplicity	0..1		
Type	EcucIntegerParamDef		
Range	0 .. 255		
Default value	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	

	Post-build time	--	
Scope / Dependency			

SWS Item	MCOS1024_Conf :		
Name	OslocInitValue		
Description	Initial Value for the data to be transferred on the IOC communication channel.		
Multiplicity	0..1		
Type	EcucStringParamDef		
Default value	--		
maxLength	--		
minLength	--		
regularExpression	--		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

SWS Item	MCOS1005_Conf :		
Name	OslocDataTypeRef		
Description	This is the type of the data to be transferred on the IOC communication channel. This attribute is necessary to generate the parameter type of the loc functions. Additionally this information should be used to compute the data size for necessary data copy operations within the loc module. If more than one attribute is defined, the IOC generator should generate an locXxxGroup function (Xxx= CHOICE [Send, Receive, Write, Read]). N:1 communication (Multiplicity of OslocSenderProperties > 1) is only allowed for multiplicity of OslocDataRef = 1		
Multiplicity	1		
Type	Foreign reference to [IMPLEMENTATION-DATA-TYPE]		
ConfigurationClass	Pre-compile time	X	All Variants
	Link time	--	
	Post-build time	--	
Scope / Dependency			

No Included Containers

10.4 Published Information

[OS766] [The standardized common published parameters as required by BSW00402 in the SRS General on Basic Software Modules [3] shall be published within the header file of this module and need to be provided in the BSW Module Description. The according module abbreviation can be found in the List of Basic Software Modules [13].] ()

Additional module-specific published parameters are listed below if applicable.

11 Generation of the OS

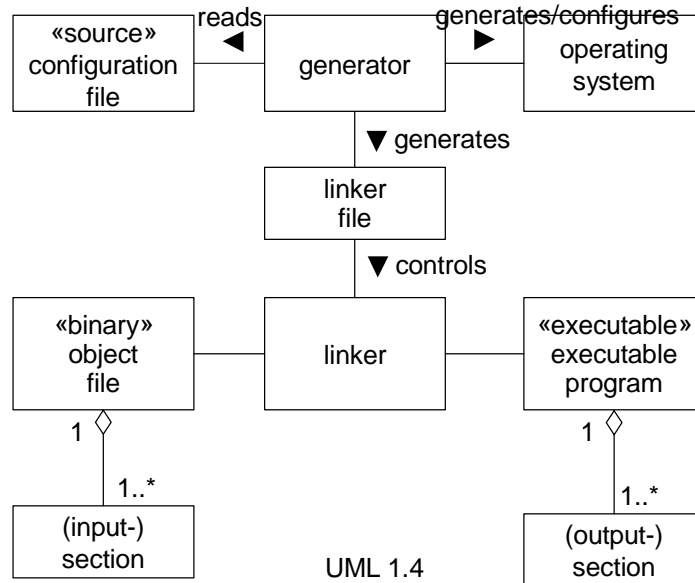


Figure 11.1: Generation Activities

11.1 Read in configuration

[OS172] [The generator shall provide the user the ability of reading the information of a selectable configuration file.] ()

11.2 Consistency check

The consistency check can issue warnings or errors. Warnings mean that the generation is completed successfully, only indicating a not advisable configuration. Errors mean that the generation is not performed.

[OS173] [The generator shall provide the user the ability of performing a consistency check of the current configuration.] ()

[OS050] [If service protection is required and `OsStatus` is not equal to `EXTENDED` (all the associated error handling is provided), the consistency check shall issue an error.] ()

[OS045] [If timing protection is configured together with OSEK OS Category 1 interrupts, the consistency check shall issue a warning.] ()

[OS562] [If timing protection is configured together with Pre- or PostTaskHook the consistency check shall issue a warning.] ()

[OS320] [If configured attributes do not match the configured scalability class (e.g. defining an execution time budget in Tasks or Category 2 ISRs and selected scalability class is 1) the consistency check shall issue a warning.] ()

[OS311] [If `OsScalabilityClass` is SC3 or SC4 AND a Task OR Category 2 ISR OR Counters OR Alarms OR Schedule tables does not belong to exactly one OS-Application the consistency check shall issue an error.] ()

[OS361] [If `OsScalabilityClass` is SC3 or SC4 AND a Category 1 ISR does not belong to exactly one trusted OS-Application the consistency check shall issue an error] ()

[OS177] [If `OsScalabilityClass` is SC3 or SC4 AND an interrupt source that is used by the OS is assigned to an OS-Application, the consistency check shall issue an error.] ()

[OS303] [If `OsAlarmIncrementCounter` is configured as action on alarm expiry AND the alarm is driven directly or indirectly (a cyclic chain of alarm actions with `OsAlarmIncrementCounter`) by that counter, the consistency check shall issue a warning..] ()

[OS328] [If `OsStatus` is STANDARD and `OsScalabilityClass` is SC3 or SC4 the consistency check shall issue an error.] ()

[OS343] [If `OsScalabilityClass` is SC3 or SC4 AND a task is referenced within a schedule table object AND the OS-Application of the schedule table has no access to the task, the consistency check shall issue an error.] ()

[OS344] [If `OsScalabilityClass` is SC3 or SC4 AND a task is referenced within an alarm object AND the OS-Application of the alarm has no access to the task, the consistency check shall issue an error.] ()

[OS440] [If a schedule table has `OsScheduleTblSyncStrategy` = IMPLICIT and the `OsCounterMaxAllowedValue+1` of the associated counter is not equal to the duration of the schedule table then the consistency check shall issue an error.] ()

[OS461] [If `OsScalabilityClass` is `SC2`, `SC3` or `SC4` AND Alarm Callbacks are configured the consistency check shall issue an error.] ()

11.3 Generating operating system

[OS179] [If the consistency check of the read-in configuration file has not run free of errors, the generator shall not generate/configure the operating system.] ()

[OS336] [The generator shall generate a relocatable memory section containing the interrupt vector table.] (BSW11019)

[OS370] [The generator shall print out information about timers used internally by the OS during generation (e.g. on console, list file).] (SWFRT00022)

[OS393] [The generator shall create conversation macros to convert counter ticks (given as argument) into real time. The format of the macro is `OS_TICKS2<Unit>_<Counter>(ticks)` whereas `<Unit>` is one of `NS` (nanoseconds), `US` (microseconds), `MS` (milliseconds) or `SEC` (seconds) and `<Counter>` is the name of the counter; E.g. `OS_TICKS2MS_MyCounter()`] (SWFRT00047)

12 Application Notes

12.1 Hooks

In OSEK OS, PreTask & PostTask Hooks run at the level of the OS with unrestricted access rights and therefore must be trusted. It is strongly recommended that these hook routines are only used during debugging and are not used in a final product.

When an OS-Application is killed the shutdown and startup hooks of the OS-Application are not called. Cleanup of OS-Application specific data can be done in the restart task.

All application-specific hook functions (startup, shutdown and error) must return (blocking or endless loops are not acceptable).

12.2 Providing Trusted Functions

Address checking shall be done before data is accessed. Special care must be taken if parameters passed by reference point to the stack space of a task or interrupt, because this address space might no longer belong to the task or interrupt when the address is used.

The following code fragment shows an example how a trusted function is called and how the checks should be done.


```
struct parameter_struct {type1 name1, type2 name2, StatusType
return_value};

/* This service is called by the user and uses a trusted function */

StatusType system_service(
    type1 parameter1,
    type2 parameter2)
{
    /* store parameters in a structure (parameter1 and parameter2) */
    struct parameter_struct local_struct;
    local_struct.name1 = parameter1;
    local_struct.name2 = parameter2;

    /* call CallTrustedFunction with appropriate index and
    * pointer to structure */
    if(CallTrustedFunction(SYSTEM_SERVICE_INDEX, &local_struct) !=
        E_OK)
        return(FUNCTION_DOES_NOT_EXIST);
    return(local_struct.return_value);
}

/* The CallTrustedFunction() service switches to the privileged
* mode. Note that the example is only a fragment! */

StatusType CallTrustedFunction(
    TrustedFunctionIndexType ix,
    TrustedFunctionParameterRefType ref)
{
    /* check for legal service index and return error if necessary */
    if(ix > MAX_SYSTEM_SERVICE)
        return(E_OS_SERVICEID);

    /* some implementation specific magic happens: the processor is
    * set to privileged mode */
    ...

    /* indirectly call target function based on the index */
    (*(system-service_list[ix]))(ix, ref);

    /* some implementation specific magic happens: the processor is
    * set to non-privileged mode */
    ...

    return(E_OK);
}
```

```
/* This part of the system service is called by
 * CallTrustedFunction() */

void TRUSTED_system_service_part2 (TrustedFunctionIndexType a,
parameter_struct *local_struct)
{
    TaskRefType task;
    type1 parameter1;
    type2 parameter2;

    if (GetTaskID(&task) != E_OK)
        task = INVALID_TASK;

    /* get parameters out of the structure (parameter1 and
     * parameter2) */
    parameter1 = local_struct.name1;
    parameter2 = local_struct.name2;

    /* check the parameters if necessary */
    /* example is for parameter1 being an address and parameter2
     * being a size */
    /* example only for system_service called from tasks */
    if(GetISRID()!=INVALID_ISR)
    {
        /* error: not callable from ISR */
        local_struct.return_value = E_OS_ACCESS;
    }
    else if(OSMEMORY_IS_WRITEABLE(CheckTaskMemoryAccess(
        task,parameter1,parameter2)))
    {
        /* system_service_part3() is now the function as it
         * would be if directly called in a non-protected
         * environment */
        local_struct.return_value =
            system_service_part3(parameter1,parameter2);
    }
    else
    {
        /* error handling */
        local_struct.return_value = E_OS_ACCESS;
    }
}
```

Note: Since the service of `CallTrustedFunction()` is very generic, it is needed to define a stub-interface which does the packing and unpacking of the arguments (as the example show). Depending on the implementation the stub interface may be (partly) generated by the generation tool.

12.3 Migration hints for OSEKtime OS users

All important OSEKtime OS features are supported in AUTOSAR OS and it should be relatively easy to port applications from OSEKtime OS to AUTOSAR OS.

However, most OSEKtime OS features are implemented slightly differently and requiring some porting effort. The following steps show how to proceed.

- Dispatcher tables can be implemented by using schedule tables provided by AUTOSAR OS. Synchronization to a global time base can be done in a similar way to OSEKtime by using the `SyncScheduleTable()` API call. A more elegant synchronization solution is also available by driving the schedule table directly from the global time source. However, the AUTOSAR OS implements priority based scheduling rather than the stack based scheduling of OSEKtime. Therefore, priorities have to be chosen for the tasks. If a given OSEKtime dispatcher table has to be converted, all tasks can be given the same priority as long as there are no task preemptions. If this cannot be guaranteed, in each case where a task could be pre-empted at a dispatch point, the pre-empting task must be allocated a strictly higher priority than the task it pre-empts. Usually, there are few preemptions in OSEKtime systems, so the priorities are easy to calculate – a simple monotonically increasing priority assignment relative to the tasks position in the schedule table should suffice in most cases. Caveat: In OSEKtime, it is theoretically possible that task A pre-empts task B at one point in the dispatcher table and task B pre-empts task A at another point (however, this is rarely used in practice). Such behaviour is not directly possible in AUTOSAR OS. It can, however, be emulated if required, either by constructing a simple state machine in the task bodies, or by adding two tasks A' and B' using the same code as tasks A and B respectively.
- Deadline monitoring is not supported by AUTOSAR OS - instead, worst-case execution time enforcement is provided. Schedulability analysis can be used to calculate whether given deadlines are met in a system of periodic tasks with given worst-case execution times.
- Reenabling of interrupts defined offline is not supported by AUTOSAR OS.
- Tasks that have precedence over interrupt service routines are not supported by AUTOSAR OS, however, this behaviour can be easily emulated by activating a low-priority task from an ISR.
- Smooth synchronization is achieved by adjusting the delay between adjacent expiry points, generalising OSEKtime OS' approach, where the synchronization of the local time to the global time is done during several dispatcher rounds by extending or shortening the last ground state of the dispatcher round.

The OSEK time specification allows dispatcher rounds to take 3 modes:

1. Synchronous
2. Asynchronous/Hard
3. Asynchronous/Smooth

Users of OSEKtime who are migrating the AUTOSAR OS can define a schedule table that has the same range/tick resolution as their global time source (with an accompanying AUTOSAR OS counter that has the same resolution as the global time) and can synthesise these modes as follows:

1. Synchronous: Define `OsScheduleTblSyncStrategy = IMPLICIT` and start using `StartScheduleTableAbs()`. Or define `OsScheduleTblSyncStrategy = EXPLICIT` and start using `StartScheduleTableSynchron()`
2. Asynchronous/Hard: Define `OsScheduleTblSyncStrategy = EXPLICIT` and specify that the final expiry point on the schedule table has a `OsScheduleTableMaxShorten = 1` and a `OsScheduleTableMaxLengthen = OsCounterMaxAllowedValue`. Start using `StartScheduleTableRel()`.
3. Asynchronous/Smooth: Define `OsScheduleTblSyncStrategy = EXPLICIT` and specify that each expiry point on the schedule table has `OsScheduleTableMaxShorten = 1` and a `OsScheduleTableMaxLengthen < OsCounterMaxAllowedValue`. Start using `StartScheduleTableRel()`.

12.4 Software Components and OS-Applications

Trusted OS-Applications can be permitted access to IO space. As software components can not be allowed direct access to the hardware, software components can not be trusted OS-Applications because this would violate this protection feature. The configuration process must ensure that this is the case.

The AUTOSAR Virtual Function Bus (VFB) specification places no restrictions on how runnables from software components are mapped to OS tasks. However, the protection mechanisms in AUTOSAR OS apply only to OS managed objects. This means that all runnables in a task:

- Are not protected from each other at runtime
- Share the same protection boundary

If runnables need to be protected they must therefore be allocated to different tasks and those tasks protected accordingly.

A simple rule can suffice:

“When allocating runnables to tasks, only allocate runnables from the same software component into the same task.”

If multiple software components from the same application are to reside on the same processor, then, assuming protection is required between applications (or parts thereof) on the same processor, this rule could be modified to relax the scope of protection to the application:

“When allocating runnables to tasks, only allocate runnables from the same application into the same task.”

If an OS-Application is killed and the restart task is activated it can not assume that the startup of the OS-Application has finished. Maybe the fault happened in the application startup hook and no task of the application was started so far.

12.5 Global Time Synchronization

The OS currently assumes that the global time synchronization is done by the user (unless implicit synchronization is used). This allows maximum flexibility regarding the time source. For synchronization with e.g. FlexRay some glue code may be necessary which transfer the information from the time source to the OS.

12.6 Working with FlexRay

Schedule tables in the AUTOSAR OS may be synchronized with a global (network) time provided by FlexRay in essentially two ways:

1. Using the FlexRay interface's services for controlling timer interrupts related to global time to provide a "hardware" counter tick source to drive the processing of a schedule table (implicit synchronization)
2. Using the FlexRay interface's service for accessing the current global time and passing this into the OS through the SyncScheduleTable() OS service call

This section looks at the second option only.

In FlexRay time is presented as a tuple of a Cycle and a MacrotickOffset within the cycle. Cycle is an 8-bit value and MacrotickOffset is a 16-bit value.

In AUTOSAR OS a schedule table is associated with an underlying counter that has a notion of ticks. It is therefore possible to synchronize with either the Cycle or the tuple of Cycle/MacrotickOffset to give the resolution of synchronization required by the application.

If Cycle only resolution is required then an OS COUNTER object should be configured to have a OsCounterMaxAllowedValue equal to the maximum number of Cycles. If Cycle/MacrotickOffset is required then an OS COUNTER object should be configured with a OsCounterMaxAllowedValue of the maximum number of Cycles multiplied by the MacrotickOffset. This provides the OS with a time base against which a ScheduleTable can be synchronized.

Synchronization between the OS and an external global time source is provided by telling the OS the global time through the SyncScheduleTable() service call. This call takes a scalar parameter of TickType so to interface this to FlexRay's representation of time a small conversion needs to be done. The following example assumes a Cycle of 255 with 65535 Macroticks per Cycle. TickType is at least 24-bits wide.

```
#define OSTIME(x) (TickType)(x);  
FrIf_GetGlobalTime(Controller, &Cycle, &Macrotick);  
SyncScheduleTable(Tbl, ((OSTIME(Cycle) << 16)+(OSTIME(Macrotick))));
```

Telling the ScheduleTable that GlobalTime can be done when the application detects that the FlexRay controller has lost synchronization with the network (by polling the controller sync status). The following code indicates how this can be used to force an associated ScheduleTable into the `SCHEDULETABLE_RUNNING` state from the `SCHEDULETABLE_RUNNING_AND_SYNCHRONOUS` state.

```
Fr_SyncStateType CurrentSyncStatus;  
if (FrIf_GetSyncState(Controller, &CurrentSyncStatus) == E_OK) {  
    if (CurrentSyncStatus == FR_ASYNC ) {  
        SetScheduleTableAsync(Table);  
    }  
}
```

Of course, other actions are possible here, like stopping the ScheduleTable, as best fits user requirements.

12.7 Migration from OIL to XML

This version of the AUTOSAR OS specification does not directly support the configuration via OIL. The support for OIL was dropped in favour of XML because XML is the standard configuration language in AUTOSAR and is essential if configuration data has to be imported / exported from / to other AUTOSAR modules or between different tools during development.

Since OIL and XML are both ASCII formats a tool vendor may offer a possibility to import (old) OIL files and to store them as (AUTOSAR OS) XML files. Currently all known vendors support at least the import of existing OIL configurations.

Note that for showing conformance to the OSEK OS specification, each OSEK OS vendor must support OIL. This means that practically each AUTOSAR OS vendor will offer some sort of import of OIL configurations – at least to show the OSEK OS conformance.

12.8 Migrating RES_SCHEDULER in AUTOSAR OS

As stated in 7.1.2.1 AUTOSAR OS treats `RES_SCHEDULER` as a normal resource. If you have legacy code which is migrated to AUTOSAR OS please take care of the following aspects:

- In OSEK OS there is no need to configure the `RES_SCHEDULER` in the OIL file. If you migrate to AUTOSAR OS the configuration is done in XML and each resource must be properly configured. The easiest way to do this is to configure a resource `RES_SCHEDULER` in XML (OsResource) and allow any Task in your system to use this resource⁷.
- Avoid that ISRs are using the `RES_SCHEDULER`. In OSEK OS this is also not possible.

⁷ This work can be done automatically by a configuration tool during importing an OIL file
222 of 230

- Make the `RES_SCHEDULER` a `STANDARD` resource (at least not an `INTERNAL` resource). The symbol `RES_SCHEDULER` must be present which is not the case if the resource is an `INTERNAL` resource.
- If you are using OS-Applications, the `RES_SCHEDULER` should belong to a trusted OS-Application. Tasks of other OS-Applications should be configured to have the right to access the resource.

12.9 Debug support

For the AUTOSAR OS the following information may be useful for users and should be considered for debug support (and may be published, e.g. in the BSWMD):

- General information about how to retrieve the current (active) Task or ISR and their (current) priority and (current) stack.
- For ISRs: Information about the name of interrupts, their mapping to the ISR identifier, the associated hardware and the used stack(s).
- For Tasks: Information about the name of the Task, its identifier, the task state, the possible priorities, the event mask (if its an extended task), the OS-Application to whom the Task belongs (if existant) and the used stack.
- For Resources: Information about the name of the Resource, its mapping to the identifier, its priority and the current owner (the Task/ISR which currently holds the Resource)
- For Alarms: Information about the name of the Alarm, its mapping to the identifier, the counter to whom it belong, the action which is executed on expiry and the current state (running or stopped). In running state the next expiry in ticks and the possible cycle time shall be also published.
- For Counters: Information about the name of the Counter, its mapping to the identifier, its associated alarms and the current counter value.
- For Schedule Tables: Information about the name of the Schedule Table, its mapping to the identifier, its current state and the next expiry point (if the table is running).
- For OS-Applications: Information about the name of the OS-Application, its mapping to the identifier, its current state and the memory sections assigned to it (if memory protection is used).

User documentation should contain information about the implemented debug features.

12.10 Integration hints for peripheral protection

Peripheral protection requires configuration on the core level usually conditioned by a supervisor access. For this reason the task of the peripheral protection is assigned to the OS module.

Peripheral protection may be implemented in two ways

- using MPU
- using dedicated peripheral protection units of the target MCU.

When using the memory protection unit, it is reasonable if two or more protected region descriptors are available for peripheral protection mechanism. The region descriptors shall be programmed to allow access to those peripherals the current OS-Application shall work with. The defined regions shall cover all memory mapped configuration registers for the peripherals to be protected. The advantage of using the MPU is that the configuration is the same as for memory protection. One of the disadvantages of this method is that it could be impossible to cover all peripheral control registers with available MPU region descriptors. The number of such descriptors is typically low.

Beware that using this method may have implication to the linker file of the project software configuration.

Second method is using a dedicated register protection schema. This method shall allow to precisely select peripherals for every OS Application. However the number of peripherals may make the register protection implementation rather bulky. Therefore it is advisable to reduce the number of protected peripherals to a reasonable value.

For both methods the configuration shall be placed into custom OS Application properties. The configuration shall be active when a task (or ISR) of a particular OS Application is running.

13 AUTOSAR Service implemented by the OS

13.1 Scope of this Chapter

This chapter is an addition to the specification of the Operating System. Whereas the other parts of the specification define the behavior and the C-interfaces of the OS module, this chapter formally specifies the corresponding AUTOSAR Service in terms of the SWC Template. The interfaces described here will be visible on the VFB and are used by the RTE generator to create the glue code between the application software (SWC) and the OS.

13.1.1 Package

The following definitions are interpreted to be in
`ARPackage AUTOSAR/Services/Os`

13.2 Overview

The AUTOSAR Operating System is normally not used directly by SWCs. Even the other BSW modules which are below the RTE are using the BSW Scheduler to have access to OS services. The BSW Scheduler of course uses the OS to implement its features, e.g. critical sections.

Nevertheless there are some cases, where it makes sense to allow SWCs access to services of the OS:

- **Timer services**
Since the number of timers in an ECU is limited it make sense to share these units across several SWCs. The functionality of the timer services of the OS which are offered to the SWCs are:
 - A service to get the current value of a – hardware or software – counter
 - A service which calculates the time difference between the current timer value and a given (previols read) timer value
- **Application modes**
An application mode is always used to start the OS. To get the current application mode the corresponding OS service is available to SWCs.
- **OS-Application handling**
To enable SWCs to start and stop OS-Applications the following services are available:
 - A service to terminate and optionally restart an OS-Application
 - A service to get the current state of the OS-Application

13.3 Specification of the Ports and Port Interfaces

This chapter specifies the ports and port interfaces which are needed in order to operate the timer services of the OS over the VFB. Note that there are ports on both sides of the RTE: The SW-C description of the OS service will define the ports below the RTE. Each SW-Component, which uses the Service, must contain “service ports” in its own SW-C description which will be connected to the ports of the OS, so that the RTE can be generated.

13.3.1 Data Types and Port Interface

13.3.1.1 General Approach

It is appropriate to model the requests issued from a client to the services by ports using the client/server interfaces.

13.3.1.2 Data Types

This chapter describes the data types which will be used in the port interfaces for service requests. In general the interfaces are using the following types:

- CounterType – This type is the reference to the requested Counter
- TickType – This type holds a timer value
- AppModeType – This type holds the current mode of the OS
- ApplicationType – This type is the reference to a OS-Application
- RestartType – This type holds the restart parameter
- ApplicationStateType – This type holds the state of an OS-Application

The APIs of the services have a return type of `StatusType`. This means that a successful call returns 0 and a return value not equal 0 represents an error.

13.3.1.3 Port Interface

The operations correspond to the function calls of the OS C-API (notation in pseudo code; must be transferred into XML).

The notation of possible error codes resulting from server calls follows the approach in the meta-model. It is a matter of the RTE specification [9], how those error codes will be passed via the actual API.

```
[OS560] [
ClientServerInterface OsService {
    PossibleErrors {
        E_OS_ACCESS = 1
        E_OS_ID = 3,
        E_OS_STATE = 7
        E_OS_VALUE = 8
    }
}
```

```

};

// The timer services

GetCounterValue(      IN CounterType CounterID,
                      OUT TickType   Value
                      ERR{E_OS_ID});

GetElapsedValue(      IN CounterType CounterID,
                      INOUT TickType PreviousValue,
                      OUT TickType   Value,
                      ERR{E_OS_ID, E_OS_VALUE});

// Service to access the current AppMode (which was
// used in StartOS()).

GetActiveApplicationMode(OUT AppModeType CurrentMode);

// Services to terminate applications and to access the current
// application state.

TerminateApplication( IN ApplicationType Application,
                     IN RestartType   RestartOption,
                     ERR{E_OS_ID, _E_OS_VALUE,
                          E_OS_STATE, E_OS_ACCESS});

GetApplicationState( IN ApplicationType Application,
                     OUT ApplicationStateType Value,
                     ERR{E_OS_ID}));

];
] ( )

```

13.3.1.4 Ports

We end up with the following structure for the AUTOSAR Interface of the OS:

```

[OS561] [
Service Os
{
    ProvidePort OsService OsService;
};

```

It is obvious that the existence of all these port definitions depends on the ECU.] ()

14 Outlook on Memory Protection Configuration

As stated before, memory protection configuration is not standardized yet. Nevertheless it seems helpful to contribute a recommendation in this chapter, how the configuration might work.

14.1 Configuration Approach

Both, SW-Components and BSW modules, map code and variables to dedicated, disjoined memory sections (see meta-class »ObjectFileSection« in chapter 7.3 of »Software Component Template«, Version 2.0.1, and »module specific sections« in chapter 8.2 of »Specification of Memory Mapping«, Version 1.0.1).

This essential precondition (avoid an inseparable conglomeration of variables in the default section) can be used to support configuration of memory protection domains:

1. The generator can save for each OS-Application a (processor-specific) maximum number of output sections for data in a file (to be used in the linker file).
2. The generator can uniquely identify the address spaces of the data output sections with symbols using the naming convention (see »memory allocation keywords« `_STOP_SEC_VAR` and `_START_SEC_VAR` for start and stop symbols) in the specification mentioned above.

The input data sections in the object files of an OS-Application can then be assigned to the output sections (with potential tool support). Usually, this is one segment for global data, and one segment for code.

To achieve portability, the user shall group all variables belonging to a private data section (Task/ISR or OS-Application) in separate files.

15 Changes to Release 3.0/3.1

- Many small correction (wording, typos, clarifications)
- Added additional services to the service interface for SWCs
- Changes caused by R4.0 concepts (e.g. debugging concept, error handling concept, multicore concept, ...)
 - o Added states to OS-Applications
 - o Added 2 new services: GetApplicationState() and AllowAccess()
 - o Extended API to terminate other OS-Applications

16 Not applicable requirements

[OS767] [These requirements are not applicable to this specification.] (BSW00344, BSW00404, BSW00405, BSW170, BSW00380, BSW00419, BSW00381, BSW00412, BSW00383, BSW00384, BSW00375, BSW00406, BSW168, BSW00407, BSW00423, BSW00337, BSW00338, BSW00369, BSW00339, BSW00422, BSW00417, BSW00409, BSW00385, BSW00386, BSW00437, BSW161, BSW162, BSW00415, BSW00325, BSW00326, BSW00342, BSW007, BSW00413, BSW00347, BSW00441, BSW00305, BSW00307, BSW00310, BSW00373, BSW00327, BSW00335, BSW00350, BSW00410, BSW00411, BSW00314, BSW00370, BSW00435, BSW00436, BSW00361, BSW00301, BSW00302, BSW00328, BSW00312, BSW006, BSW00439, BSW00357, BSW00377, BSW00304, BSW00355, BSW00378, BSW00306, BSW00308, BSW00309, BSW00358, BSW00414, BSW00376, BSW00440, BSW00329, BSW00330, BSW009, BSW00401, BSW172, BSW010, BSW00333, BSW00374, BSW00379, BSW003, BSW00318, BSW00321, BSW00341, BSW00334, SWFRT00032)