

<b>Document Title</b>	Explanation of Adaptive Platform Software Architectural Decisions
<b>Document Owner</b>	AUTOSAR
<b>Document Responsibility</b>	AUTOSAR
<b>Document Identification No</b>	983

<b>Document Status</b>	published
<b>Part of AUTOSAR Standard</b>	Adaptive Platform
<b>Part of Standard Release</b>	R21-11

<b>Document Change History</b>			
<b>Date</b>	<b>Release</b>	<b>Changed by</b>	<b>Description</b>
2021-11-25	R21-11	AUTOSAR Release Management	<ul style="list-style-type: none"> <li>Added architectural decisions made for release R21-11</li> <li>Updated the list of affected functional clusters</li> </ul>
2020-11-30	R20-11	AUTOSAR Release Management	<ul style="list-style-type: none"> <li>Initial release</li> </ul>

## Disclaimer

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

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

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

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

The word AUTOSAR and the AUTOSAR logo are registered trademarks.

## Table of Contents

1	Introduction	4
1.1	Objectives	4
1.2	Scope	4
2	Definition of Terms and Acronyms	5
2.1	Acronyms and Abbreviations	5
2.2	Definition of Terms	5
3	Related Documentation	6
4	Overview	7
5	Architectural Decisions	8
5.1	Dynamic memory allocation	8
5.2	Making Adaptive Runtime classes final	9
5.3	Usage of out parameters	10
5.4	Usage of named constructors for exception-less object creation	11
5.5	Introduction of a monotonic clock API	12
5.6	Responsibilities of State Management, Execution Management, and Platform Health Management	13
5.7	Use of local proxy objects for shared access to objects	16
5.8	Harmonized production errors	17
5.9	Default arguments are not allowed in virtual functions	18
5.10	Assert that only APIs from properly initialized functional clusters can be called	19
5.11	Provide only interfaces that are intended to be used by AUTOSAR applications and other Functional Clusters	20
5.12	AUTOSAR should follow the C++ Core Guidelines	21
5.13	Harmonized error handling for lost daemon connection	22
5.14	Granularity of diagnostics	23
5.15	Faults inside constructor	24
5.16	Restart granularity	26
5.17	Platform-independent development of Software Clusters of category APPLICATION_LAYER	26
5.18	Harmonized logging / tracing	27

# 1 Introduction

This explanatory document provides additional information on architectural decisions made for the AUTOSAR Adaptive Platform standard.

## 1.1 Objectives

The main objective of this document is to provide a **documentation of architectural decisions** made for the AUTOSAR Adaptive Platform to make such decisions comprehensible and reviewable in the future and ultimately get a more maintainable standard.

## 1.2 Scope

This document covers decisions made for the software architecture of the AUTOSAR Adaptive Platform. The main audience of this document are architects of the AUTOSAR Adaptive Platform as well as members of other working groups.

## 2 Definition of Terms and Acronyms

### 2.1 Acronyms and Abbreviations

Abbreviation / Acronym	Description
API	Application Programming Interface
STL	Standard Template Library

### 2.2 Definition of Terms

Term	Description
Adaptive Application	See [1, AUTOSAR Glossary].
Execution Management	A <code>Functional Cluster</code> . See [2, EXP_SWArchitecture] for an overview.
Functional Cluster	A logical group of functionality within the AUTOSAR Adaptive Platform. Functional Clusters are the subject of the individual specification documents that make up the AUTOSAR Adaptive Platform standard. See [2, EXP_SWArchitecture] for an overview of all Functional Clusters in the AUTOSAR Adaptive Platform.
Platform Health Management	A <code>Functional Cluster</code> . See [2, EXP_SWArchitecture] for an overview.
Process	See [1, AUTOSAR Glossary].
State Management	A <code>Functional Cluster</code> . See [2, EXP_SWArchitecture] for an overview.
Software Cluster	See [1, AUTOSAR Glossary] and [2, EXP_SWArchitecture].
Thread	See [1, AUTOSAR Glossary].
Watchdog	An external component that supervises execution of the AUTOSAR Adaptive Platform. See [2, EXP_SWArchitecture] for an overview.

### 3 Related Documentation

This document provides an overview of the architectural decisions that have been made for the AUTOSAR Adaptive Platform and their rationale. A high-level overview of the current AUTOSAR Adaptive Platform architecture is provided in [[2](#), EXP\_SWArchitecture].

## 4 Overview

This chapter provides an overview of the organization and structure of decisions listed in this document. All decisions are organized in a table as described below.

<b>Date of approval</b>	Date when the decision has been approved by CF-CCB.
<b>Decision</b>	The decision itself. Impact or direct consequences (for example, changes to interfaces) of the decision are not documented in the normal case. Changes are requested during the roll-out process after the decision has been made.
<b>Rationale</b>	A rationale for the decision.
<b>Architectural cluster</b>	A list of architectural clusters that are affected by the decision.
<b>Category</b>	Category of the decision.
<b>Application affected</b>	States if the decision has a direct impact on existing applications.
<b>Assumptions</b>	Lists the assumptions that have been made before making the decision itself. These assumptions are documented in order to be able to review decisions in the future and check if some assumptions probably no longer hold.
<b>Constraints</b>	Provides an overview of the constraints that were identified to have an impact on possible solutions. The constraints are also documented in order to be reference points for future reviews of the decision.
<b>Alternatives</b>	Lists the alternatives that were considered and a rationale why they are worse than the decision that has been made.
<b>Remarks</b>	Lists remarks on the decisions.
<b>Related requirements</b>	Lists requirements related to the decision.

## 5 Architectural Decisions

This chapter lists architectural decisions that have been made for the AUTOSAR Adaptive Platform.

### 5.1 Dynamic memory allocation

<b>Date of approval</b>	2020-09-15
<b>Decision</b>	The use of dynamic memory allocation by Adaptive Applications and Functional Clusters is allowed and assumed upon designing the AUTOSAR Adaptive Platform standard.
<b>Rationale</b>	<p>The use of dynamic memory allocation is essentially indispensable as the AUTOSAR Adaptive Platform standard employs C++ as the language for its API.</p> <p>As the AUTOSAR Adaptive Platform standard will be used in safety-related systems, dynamic memory allocation can cause non-deterministic behavior. Two typical issues are the fragmentation and non-deterministic allocation/de-allocation processing time. Memory allocators designed for non-safety-critical systems often exhibit such issues, as they are more or less designed for memory efficiency and/or average processing performance.</p> <p>These issues can be controlled by using deterministic memory allocators. Memory allocation is a well-studied area and various techniques have been reported (Refer to references below for some examples). Multiple AUTOSAR partners within the architecture group reportedly have such deterministic memory allocators implemented and have been used in mass-production systems.</p> <p>Note that such allocators should replace the default <code>malloc()/free()</code> implementations provided in the standard C library, that sits underneath the C++ runtime library providing <code>new()/delete()</code> and also STL that AUTOSAR Adaptive Platform also uses. This frees applications from providing its own custom deterministic allocators and installing it to custom-allocator-aware classes.</p> <p>Please refer to [3], [4], [5], and [6] for further information on memory fragmentation and memory allocation in real-time systems.</p>
<b>Architectural cluster</b>	<ul style="list-style-type: none"> <li>• Communication Management</li> <li>• Adaptive Core</li> <li>• Intrusion Detection System Manager for Adaptive Platform</li> <li>• Persistency</li> <li>• Execution Management</li> <li>• Time Synchronization</li> <li>• Platform Health Management</li> </ul>







	<ul style="list-style-type: none"> <li>• Identity and Access Management</li> <li>• Diagnostic Management</li> <li>• Log and Trace</li> <li>• Cryptography</li> <li>• State Management</li> <li>• Update and Configuration Management</li> <li>• Network Management</li> </ul>
<b>Category</b>	Safety
<b>Application affected</b>	No
<b>Assumptions</b>	Platform vendors and/or compiler vendors can replace the default memory allocation/deallocation functions to use deterministic versions of those functions during critical phases of the runtime when such determinism is required for safety purposes.
<b>Constraints</b>	During certain phases of the runtime determinism is required. These are the phases in which the allocators need to be replaced with deterministic versions.
<b>Alternatives</b>	<ul style="list-style-type: none"> <li>• Not using dynamic memory allocation is not an alternative for using C++.</li> <li>• Disallow dynamic memory allocation during certain phases of the runtime in which determinism is required. This makes it very difficult to run complex code during these phases.</li> </ul>
<b>Remarks</b>	No remarks.
<b>Related requirements</b>	<ul style="list-style-type: none"> <li>• [RS_AP_00129] Public types defined by functional clusters shall be designed to allow implementation without dynamic memory allocation</li> </ul>

## 5.2 Making Adaptive Runtime classes final

<b>Date of approval</b>	2020-09-15
<b>Decision</b>	Adaptive Runtime types shall use the <code>final</code> specifier unless they are meant to be used as a base class.
<b>Rationale</b>	<p>Making classes <code>final</code> that are not intended to be used as base class expresses the design (in particular the class hierarchy) more explicit. This will avoid problems such as</p> <ul style="list-style-type: none"> <li>• to derive from a class that is not prepared for sub-classing,</li> <li>• to inadvertently create a new virtual function instead of overwriting a function from the base class due to a slightly different signature.</li> </ul>
<b>Architectural cluster</b>	<ul style="list-style-type: none"> <li>• Communication Management</li> <li>• Adaptive Core</li> <li>• Intrusion Detection System Manager for Adaptive Platform</li> </ul>





	<ul style="list-style-type: none"> <li>• Persistency</li> <li>• Execution Management</li> <li>• Time Synchronization</li> <li>• Platform Health Management</li> <li>• Diagnostic Management</li> <li>• Log and Trace</li> <li>• Cryptography</li> </ul>
<b>Category</b>	None
<b>Application affected</b>	No
<b>Assumptions</b>	A clear expression of the intended design of the public AUTOSAR Runtime for Adaptive Applications class hierarchy.
<b>Constraints</b>	No constraints were identified.
<b>Alternatives</b>	The alternative is to have a code review of the application code using AUTOSAR types. This is far out-of-scope of AUTOSAR wherefore it is not a real alternative.
<b>Remarks</b>	No remarks.
<b>Related requirements</b>	<ul style="list-style-type: none"> <li>• [RS_AP_00140] Usage of "final specifier" in ara types</li> </ul>

### 5.3 Usage of out parameters

<b>Date of approval</b>	2020-09-15
<b>Decision</b>	Out parameters can be used for in-place modifications but shall not be used for returning values.
<b>Rationale</b>	<p>Harmonized look and feel.</p> <p>C++ Core Guidelines [7]: "F.20: For "out" output values, prefer return values to output parameters. [...] A return value is self-documenting, whereas a &amp; could be either in-out or out-only and is liable to be misused. This includes large objects like standard containers that use implicit move operations for performance and to avoid explicit memory management."</p>
<b>Architectural cluster</b>	<ul style="list-style-type: none"> <li>• Persistency</li> </ul>
<b>Category</b>	None
<b>Application affected</b>	No
<b>Assumptions</b>	Dynamic memory allocation is allowed for all cases in which the APIs are used, even when running time critical safety related code including ASIL D.
<b>Constraints</b>	In/out parameters, i.e. modifying an already existing parameter within a function is allowed. For example, a function that clears or writes to a buffer should receive that buffer as a non-const in/out parameter.





<b>Alternatives</b>	No alternatives were considered.
<b>Remarks</b>	No remarks.
<b>Related requirements</b>	<ul style="list-style-type: none"> <li>• [RS_AP_00141] Usage of out parameters</li> </ul>

## 5.4 Usage of named constructors for exception-less object creation

<b>Date of approval</b>	2020-10-06
<b>Decision</b>	Exceptionless functions for creation of objects which returns an <code>ara::core::Result</code> should use the "named constructor idiom".
<b>Rationale</b>	<p>Disadvantages of constructor token approach are avoided as follows:</p> <ul style="list-style-type: none"> <li>• The constructor token type is an implementation detail of a <code>Class</code> and should thus not be specified, or even accessible from outside. This makes the use of <code>auto</code> for obtaining a token mandatory because the token type cannot be referred to in any other way.</li> <li>• Moving the token's content to the <code>SomeClass</code> instance has to be done very carefully to fulfill the always-successful guarantee, which can be tricky if multiple resources need to be acquired.</li> <li>• The token object is "destroyed" by <code>std::move</code>-ing its value into the <code>SomeClass</code> constructor (actually, it is to be in a "valid" but unspecific state according to the C++ standard), but it is easily possible to mistakenly use it again for attempting to create another instance, with undefined results.</li> </ul>
<b>Architectural cluster</b>	<ul style="list-style-type: none"> <li>• Communication Management</li> <li>• Adaptive Core</li> <li>• Intrusion Detection System Manager for Adaptive Platform</li> <li>• Persistency</li> <li>• Execution Management</li> <li>• Time Synchronization</li> <li>• Platform Health Management</li> <li>• Identity and Access Management</li> <li>• Diagnostic Management</li> <li>• Log and Trace</li> <li>• Cryptography</li> <li>• State Management</li> <li>• Update and Configuration Management</li> <li>• Network Management</li> </ul>





<b>Category</b>	Safety
<b>Application affected</b>	Yes
<b>Assumptions</b>	No assumptions were made.
<b>Constraints</b>	No constraints were identified.
<b>Alternatives</b>	<ul style="list-style-type: none"> <li>• Constructor token approach. It was not considered due to the drawbacks described in the rationale of this decision.</li> <li>• Regular constructor calls. Regular constructor calls were not considered because regular constructors may throw exceptions and thus cannot be used in an exception-less design.</li> </ul>
<b>Remarks</b>	No remarks.
<b>Related requirements</b>	No related requirements.

## 5.5 Introduction of a monotonic clock API

<b>Date of approval</b>	2020-10-20
<b>Decision</b>	<p>The AUTOSAR Runtime for Adaptive Applications shall provide its own monotonic <code>std::chrono::SteadyClock</code> representing the power-up time of the machine. The accuracy of this clock is defined by the platform vendor.</p> <p>The accuracy of this clock could be used as a characteristic value of the platform so that the projects could check whether this clock meets the project-specific requirements (e.g. time synchronization requires typically a clock with higher accuracy).</p> <p>The system start of the machine defines the epoch of the clock. So this clock represents the power-up time of the machine.</p> <p>Functional Clusters dealing with timestamps or clocks should use this clock as a basis.</p>
<b>Rationale</b>	The timestamps used in the time synchronization cluster should be based on <code>std::chrono</code> . Time synchronization requires a monotonic clock with special accuracy.
<b>Architectural cluster</b>	<ul style="list-style-type: none"> <li>• Adaptive Core</li> <li>• Execution Management</li> <li>• Time Synchronization</li> </ul>
<b>Category</b>	None
<b>Application affected</b>	Yes
<b>Assumptions</b>	The time synchronization cluster is typically a daemon-based architecture due to a single communication endpoint of the time sync messages. A standardized clock with a special accuracy as a common basis is required to synchronize the daemon with the library.
<b>Constraints</b>	No constraints were identified.





<b>Alternatives</b>	The used clock could also be passed as a template argument. But a standardized clock with a special accuracy as a common basis is required anyway in case the time synchronization cluster is daemon based.
<b>Remarks</b>	No remarks.
<b>Related requirements</b>	No related requirements.

## 5.6 Responsibilities of State Management, Execution Management, and Platform Health Management

<b>Date of approval</b>	2020-10-27
<b>Decision</b>	<p>State Management, Execution Management, and Platform Health Management are the fundament/basis of the AUTOSAR Adaptive Platform. A failure in either State Management, Platform Health Management, or Execution Management process will typically lead to stop triggering the watchdog. Platform Health Management supervises State Management and Execution Management. Platform Health Management controls the watchdog and is thus in turn supervised by the hardware watchdog.</p> <p>Triggering of a Machine reset as a last resort should not be an option at all in case of a failing of an Adaptive Application supervision (i.e. apart from Operating System/Execution Management/State Management/Platform Health Management). A supervision failure in an Adaptive Application shall be reported to State Management. State Management may forward this failure based on the criticality to Platform Health Management to wrongly trigger or stop triggering the serviced watchdog.</p> <p>Platform Health Management performs a logical supervision of checkpoints within a process or between processes within a Function Group. Platform Health Management reports any supervision failures to State Management. State Management is responsible to perform recovery actions including a switch of the Function Group State, by delegating to the Adaptive Application, or, as a last resort, by advising Platform Health Management to perform a hardware reset. Platform Health Management is intended for supervision of safety-critical processes. Thus, Platform Health Management is an optional part of the AUTOSAR Adaptive Platform for non safety-critical applications.</p> <p>Processes shall never be restarted on their own because they may have unknown runtime dependencies. The relation between a Process and a Function Group is comparable to the relation between a thread and a process. State Management should always trigger a request (Function Group State change) to restart processes even in the simplistic/non-dependent cases. Thus, Platform Health Management does not have a direct interface to Execution Management.</p> <p>The unrecoverable state interface of Platform Health Management shall be removed.</p>





<b>Rationale</b>	<p>The chosen solution leads to a simpler design of Platform Health Management with a single and well-defined responsibility. The chosen solution also adheres to the single responsibility principle for State Management (control system state) and Execution Management (control processes) as well.</p> <p>Recovery actions can be added by extension (open-closed principle) to State Management. There is no need to modify or configure Platform Health Management.</p> <p>Supervision failures may be handled by an Adaptive Application as well if State Management chooses to delegate recovery to the Adaptive Application.</p>
<b>Architectural cluster</b>	<ul style="list-style-type: none"> <li>• Execution Management</li> <li>• Platform Health Management</li> <li>• State Management</li> </ul>
<b>Category</b>	Safety
<b>Application affected</b>	Yes
<b>Assumptions</b>	<ul style="list-style-type: none"> <li>• State Management is a mandatory part of the AUTOSAR Adaptive Platform.</li> <li>• Performance impact / delay of indirect reporting of supervision failures to an Adaptive Application via State Management is negligible in comparison to execution of reasonable recovery actions (such as starting processes).</li> </ul>
<b>Constraints</b>	No constraints were identified.
<b>Alternatives</b>	<ul style="list-style-type: none"> <li>• Alternative 1 - Failure recovery coordinated by Platform Health Management</li> </ul> <p>Recovery in case of a systematic failure is coordinated by Platform Health Management. Several components (Adaptive Application, Execution Management, State Management, watchdog) are involved based on priorities. Platform Health Management coordinates the recovery in the following manner:</p> <ol style="list-style-type: none"> <li>1. Platform Health Management asks the Adaptive Application to recover</li> <li>2. In case of failure, Platform Health Management asks Execution Management to restart failed processes</li> <li>3. In case of failure, Platform Health Management asks State Management to recover by switching the Function Group State</li> <li>4. In case of failure, Platform Health Management stops triggering the watchdog and resets the Machine</li> <li>5. In case of failure, Platform Health Management switches to unrecoverable state (not yet fully defined)</li> </ol> <p>Alternative 1 was not considered due to not adhering to the single responsibility principle because several components are responsible for recovery actions. This solution would also require Platform Health Management to have application knowledge because it has to determine</p>





	<p style="text-align: center;">△</p> <p>the appropriate Function Group State in step 3. Restarting single processes may not be appropriate (step 2) due to runtime dependencies.</p> <ul style="list-style-type: none"> <li>Alternative 2 - Distributed failure recovery</li> </ul> <p>Recovery in case of a systematic failure is coordinated by Platform Health Management and State Management. Several components (Adaptive Application, Execution Management, watchdog) are involved based on priorities. Platform Health Management and State Management coordinate the recovery in the following manner:</p> <ol style="list-style-type: none"> <li>Platform Health Management asks the Adaptive Application to recover</li> <li>In case of failure, Platform Health Management asks State Management to coordinate recovery by restarting the application</li> <li>State Management asks Execution Management to change state / switch to degraded state or safe state</li> <li>In case of failure, State Management asks Adaptive Application to recover</li> <li>In case step 2 failed due to application dependencies, Platform Health Management stops triggering the watchdog and resets the Machine</li> </ol> <p>Alternative 2 was not considered due to not adhering to the single responsibility principle because Platform Health Management and State Management share responsibility for coordinating recovery actions.</p>
<p><b>Remarks</b></p>	<ul style="list-style-type: none"> <li>According to ISO 26262, it has to be ensured that a reaction is triggered after a safety-relevant failure occurred. Therefore, Platform Health Management shall make sure that State Management receives the notification on a detected failure even if they communicate via an unreliable communication channel, for example, an inter-process communication mechanism. To achieve this, Platform Health Management should implement a timeout monitoring. If no response by State Management is received after a configurable timeout and number of tries, Platform Health Management shall trigger a reaction via hardware Watchdog.</li> <li>For release R19-11 of the AUTOSAR Adaptive Platform, the configuration of Platform Health Management included rules for monitoring (PhmSupervision), arbitration and recovery actions. With this decision, Platform Health Management is only responsible for monitoring. The rules for monitoring (PhmSupervision) are unaffected. However, the responsibilities for arbitration and recovery actions are moved to State Management. In the current design, State Management is a piece of project-specific, coded software with only little configuration. The configuration for State Management should be extended to support arbitration and recovery actions as well. This will allow to validate such configurations based on standardized rules which is extremely hard to achieve on source code level.</li> </ul>
<p><b>Related requirements</b></p>	<p>No related requirements.</p>

## 5.7 Use of local proxy objects for shared access to objects

<b>Date of approval</b>	2020-10-27
<b>Decision</b>	Local proxy object(s) shall be used to provide shared access to object instance(s) via the AUTOSAR Runtime for Adaptive Applications interface.
<b>Rationale</b>	<p>Local proxy objects hide the implementation details of the shared access. The AUTOSAR Runtime for Adaptive Applications interface shall return a proxy object by value. The caller shall use the object as a local proxy for subsequent communication. Return by value is the most straightforward way to return data. This decision enforces harmonization of the AUTOSAR Runtime for Adaptive Applications interface. Stack vendors may freely choose how to implement the shared access inside the proxy class.</p> <p>An example for the use of a local proxy object by the caller is the following:</p> <pre> 1 2 Result&lt;void&gt; myFunc() { 3     Result&lt;KeyValueStorage&gt; kvsRes 4         = KeyValueStorage::Create(KVS_ID); 5 6     if (kvsRes) { 7         KeyValueStorage kvs = std::move(kvsRes).Value(); 8         auto keyRes = kvs.GetAllKeys(); // Value semantics 9         // ... 10    } else { 11        return {std::move(kvsRes).Error()}; 12    } </pre>
<b>Architectural cluster</b>	<ul style="list-style-type: none"> <li>• Communication Management</li> <li>• Adaptive Core</li> <li>• Intrusion Detection System Manager for Adaptive Platform</li> <li>• Persistency</li> <li>• Execution Management</li> <li>• Time Synchronization</li> <li>• Platform Health Management</li> <li>• Identity and Access Management</li> <li>• Diagnostic Management</li> <li>• Log and Trace</li> <li>• Cryptography</li> <li>• State Management</li> <li>• Update and Configuration Management</li> <li>• Network Management</li> </ul>
<b>Category</b>	None







<b>Application affected</b>	Yes
<b>Assumptions</b>	No assumptions were made.
<b>Constraints</b>	No constraints were identified.
<b>Alternatives</b>	The alternative of using proxy classes is the usage of handles. These handles would however reveal the implementation details of the shared access.
<b>Remarks</b>	No remarks.
<b>Related requirements</b>	<ul style="list-style-type: none"> <li>• [RS_AP_00135] Avoidance of shared ownership</li> </ul>

## 5.8 Harmonized production errors

<b>Date of approval</b>	2021-04-20
<b>Decision</b>	Functional clusters shall standardize production errors for common use-cases demanded by the market. The standardization shall summarize all production errors by a standardized table in all SWS documents specifying production errors.
<b>Rationale</b>	Production errors are a fact. In order to be able to (semi-)automatically analyze them and react to them, they and their documentation/persistence and their healing needs to be standardized.
<b>Architectural cluster</b>	<ul style="list-style-type: none"> <li>• Communication Management</li> <li>• Adaptive Core</li> <li>• Intrusion Detection System Manager for Adaptive Platform</li> <li>• Persistency</li> <li>• Execution Management</li> <li>• Time Synchronization</li> <li>• Platform Health Management</li> <li>• Identity and Access Management</li> <li>• Diagnostic Management</li> <li>• Log and Trace</li> <li>• Cryptography</li> <li>• State Management</li> <li>• Update and Configuration Management</li> <li>• Network Management</li> </ul>
<b>Category</b>	None
<b>Application affected</b>	Yes





<b>Assumptions</b>	Conceptually production errors are taken over from the AUTOSAR Classic Platform. A differentiation between production errors and extended production errors is not necessary.
<b>Constraints</b>	No constraints were identified.
<b>Alternatives</b>	Functional clusters provide interfaces to allow applications to monitor production errors.
<b>Remarks</b>	None
<b>Related requirements</b>	None

## 5.9 Default arguments are not allowed in virtual functions

<b>Date of approval</b>	2021-06-01
<b>Decision</b>	Default arguments shall not be used at all in virtual functions.
<b>Rationale</b>	The according RQ of the "C++ core guidelines" are too weak .. (they state, that it needs be made sure that a default argument is always the same) ... this would lead to code duplication with dependencies and high risks of inconsistencies, which can easily lead to unexpected behavior.
<b>Architectural cluster</b>	<ul style="list-style-type: none"> <li>• Communication Management</li> <li>• Adaptive Core</li> <li>• Intrusion Detection System Manager for Adaptive Platform</li> <li>• Execution Management</li> <li>• Time Synchronization</li> <li>• Platform Health Management</li> <li>• Diagnostic Management</li> <li>• Log and Trace</li> <li>• Cryptography</li> </ul>
<b>Category</b>	None
<b>Application affected</b>	Yes
<b>Assumptions</b>	No assumptions were made.
<b>Constraints</b>	No constraints were identified.
<b>Alternatives</b>	No alternatives were considered.
<b>Remarks</b>	No remarks.
<b>Related requirements</b>	<ul style="list-style-type: none"> <li>• [RS_AP_00148] Default arguments are not allowed in virtual functions</li> </ul>

## 5.10 Assert that only APIs from properly initialized functional clusters can be called

<b>Date of approval</b>	2021-06-08
<b>Decision</b>	If functionality is called that depends on prior initialization via <code>ara::core::Initialize</code> and <code>ara::core::Initialize</code> has not been called, the functional cluster implementation shall treat this as a violation and shall follow <code>SWS_CORE_00003</code> from [8, Specification of Adaptive Platform Core].
<b>Rationale</b>	Calling APIs from uninitialized functional clusters that depend on prior initialization cannot perform properly. This results in undefined behavior. The problem is typically caused by misconfiguration or incomplete initialization at an earlier stage of the system startup. This cannot be handled by the caller of the API at the point in time where the error is detected. Aborting execution is the only way to signal this kind of systematic error and prevent later failures.
<b>Architectural cluster</b>	<ul style="list-style-type: none"> <li>• <code>Communication Management</code></li> <li>• <code>Adaptive Core</code></li> <li>• <code>Intrusion Detection System Manager for Adaptive Platform</code></li> <li>• <code>Persistency</code></li> <li>• <code>Execution Management</code></li> <li>• <code>Time Synchronization</code></li> <li>• <code>Platform Health Management</code></li> <li>• <code>Identity and Access Management</code></li> <li>• <code>Diagnostic Management</code></li> <li>• <code>Log and Trace</code></li> <li>• <code>Cryptography</code></li> <li>• <code>State Management</code></li> <li>• <code>Update and Configuration Management</code></li> <li>• <code>Network Management</code></li> </ul>
<b>Category</b>	None
<b>Application affected</b>	Yes
<b>Assumptions</b>	Parts of the system need to be initialized statically.
<b>Constraints</b>	No constraints were identified.
<b>Alternatives</b>	<ul style="list-style-type: none"> <li>• Extend every API that depends on prior initialization with a specific error code (e.g. <code>kNotInitialized</code>) and force callers to check this error code at every call (and let them abort themselves).</li> </ul>
<b>Remarks</b>	No remarks.
<b>Related requirements</b>	None

## 5.11 Provide only interfaces that are intended to be used by AUTOSAR applications and other Functional Clusters

<b>Date of approval</b>	2021-06-08
<b>Decision</b>	<p>It is explicitly prohibited to standardize implementation details, like:</p> <ul style="list-style-type: none"> <li>• Classes, base-classes, functions etc. that are not used on the application level or in platform extension APIs</li> <li>• Implementation inheritance in the public APIs</li> <li>• C++ SFINAE techniques of any kind</li> <li>• Private members of classes</li> </ul>
<b>Rationale</b>	<ul style="list-style-type: none"> <li>• Provide only narrow interfaces to avoid coupling to implementation details.</li> <li>• Hide implementation details because by AUTOSAR definition the implementation details are on the platform vendor.</li> </ul>
<b>Architectural cluster</b>	<ul style="list-style-type: none"> <li>• Communication Management</li> <li>• Adaptive Core</li> <li>• Intrusion Detection System Manager for Adaptive Platform</li> <li>• Persistency</li> <li>• Execution Management</li> <li>• Time Synchronization</li> <li>• Platform Health Management</li> <li>• Identity and Access Management</li> <li>• Diagnostic Management</li> <li>• Log and Trace</li> <li>• Cryptography</li> <li>• State Management</li> <li>• Update and Configuration Management</li> <li>• Network Management</li> </ul>
<b>Category</b>	None
<b>Application affected</b>	Yes
<b>Assumptions</b>	No assumptions were made.
<b>Constraints</b>	No constraints were identified.
<b>Alternatives</b>	No alternatives were considered.
<b>Remarks</b>	No remarks.
<b>Related requirements</b>	<ul style="list-style-type: none"> <li>• [RS_AP_00150] Provide only interfaces that are intended to be used by AUTOSAR applications and other Functional Clusters</li> </ul>

## 5.12 AUTOSAR should follow the C++ Core Guidelines

<b>Date of approval</b>	2021-06-08
<b>Decision</b>	AUTOSAR C++ APIs should follow the [7, C++ Core Guidelines]. The exceptions for hard-real-time systems shall apply. The AUTOSAR guidelines defined in RS-General shall overrule the "C++ Core Guidelines" in case of conflict. If a part of the AUTOSAR C++ API cannot follow the "C++ Core Guidelines" for some other reason, its specification shall state the rationale (how this is done in detail, shall be aligned with the architecture group).
<b>Rationale</b>	These guidelines are well accepted in the market. Their aim is to help C++ programmers writing simpler, more efficient, and more maintainable code. Specific guidelines for the automotive domain for C++ 14 are not available. When the upcoming version of the MISRA C++ standard is published, this decision/requirement may be replaced by a decision/requirement to follow MISRA C++.
<b>Architectural cluster</b>	<ul style="list-style-type: none"> <li>• Communication Management</li> <li>• Adaptive Core</li> <li>• Intrusion Detection System Manager for Adaptive Platform</li> <li>• Persistency</li> <li>• Execution Management</li> <li>• Time Synchronization</li> <li>• Platform Health Management</li> <li>• Identity and Access Management</li> <li>• Diagnostic Management</li> <li>• Log and Trace</li> <li>• Cryptography</li> <li>• State Management</li> <li>• Update and Configuration Management</li> <li>• Network Management</li> </ul>
<b>Category</b>	None
<b>Application affected</b>	Yes
<b>Assumptions</b>	No assumptions were made.
<b>Constraints</b>	Some exceptions apply like the exception-less handling of the ARA APIs.
<b>Alternatives</b>	No alternatives were considered.
<b>Remarks</b>	No remarks.
<b>Related requirements</b>	<ul style="list-style-type: none"> <li>• [RS_AP_00151] C++ Core Guidelines</li> </ul>

## 5.13 Harmonized error handling for lost daemon connection

<b>Date of approval</b>	2021-06-08
<b>Decision</b>	<p>If a functional cluster communicates with a remote peer (e.g. IPC communication to a daemon) adequate error cases for communication failures shall be identified (e.g. lost communication). These error cases shall be grouped (according to the same error recovery mechanism) and if the user of the API shall receive notification (e.g. by callbacks or returning error codes) for a particular group, a suitable notification mechanism shall be selected.</p> <p>If notification of client is required as immediate action on error occurrence the notification mechanism shall be based on client callback. This mechanism uses registration of a state change callback handler before a client can make use of a service.</p> <p>If notification shall take place on communication attempt one of the following options shall be implemented:</p> <ul style="list-style-type: none"> <li>• provision of error code, e.g. <code>kServiceNotAvailable</code> of type <code>ara::core::ErrorDomain::CodeType</code>.</li> <li>• mapping to functional status information inside a data structure (e.g. class object), which represent an error status</li> </ul>
<b>Rationale</b>	The provider of the service call shall decide individually the most suitable mechanism if it is useful to inform application about lost daemon.
<b>Architectural cluster</b>	<ul style="list-style-type: none"> <li>• Communication Management</li> <li>• Adaptive Core</li> <li>• Intrusion Detection System Manager for Adaptive Platform</li> <li>• Persistency</li> <li>• Execution Management</li> <li>• Time Synchronization</li> <li>• Platform Health Management</li> <li>• Identity and Access Management</li> <li>• Diagnostic Management</li> <li>• Log and Trace</li> <li>• Cryptography</li> <li>• State Management</li> <li>• Update and Configuration Management</li> <li>• Network Management</li> </ul>
<b>Category</b>	None
<b>Application affected</b>	Yes





<b>Assumptions</b>	<p>The following assumptions were made:</p> <ul style="list-style-type: none"> <li>• The implementation does not depend on the type of communication interface, e.g. process local, ara::com or native IPC mechanisms are in scope of the decision.</li> <li>• There is no polling of communication status required by user of the API.</li> <li>• The cause of disconnected service shall be kept agnostic to the user of the API.</li> <li>• Connection oriented communication is out of scope due to inherent detection mechanisms of the protocol.</li> </ul>
<b>Constraints</b>	No constraints were identified.
<b>Alternatives</b>	No alternatives were considered.
<b>Remarks</b>	No remarks.
<b>Related requirements</b>	None

## 5.14 Granularity of diagnostics

<b>Date of approval</b>	2021-06-08
<b>Decision</b>	Diagnostic entity shall be identical to the deployable unit within a vehicle. Deployable unit means either hardware units (ECUs), partitions or Root-SoftwareClusters.
<b>Rationale</b>	<p>AUTOSAR focused on the SoftwareCluster approach because it offers a more easy option to keep the two worlds consistent. The Root-SoftwareCluster is the individual deployable unit from the OEM perspective. Wherefore it is easy to keep the offboard world consistent if the diagnostic has identical boundaries.</p> <p>The production and workshop systems are often bound to the physical device wherefore many OEMs want to start also with this approach in Adaptive. Consequently, until there is no individual software setup with a car (e.g. because the installed options can be chosen by the driver itself) the offboard systems could be kept consistent by stringent workflows.</p>
<b>Architectural cluster</b>	<ul style="list-style-type: none"> <li>• Diagnostic Management</li> </ul>
<b>Category</b>	None
<b>Application affected</b>	No
<b>Assumptions</b>	DM core doesn't mind if a further diagnostic server is installed (in the context of a new SoftwareCluster) or the current diagnostic server is just extended. Partitions could be Classic Platform, Adaptive Platform or non-AUTOSAR ones.





<b>Constraints</b>	Diagnostics is a (non-verbose) offboard-communication using external description to document the communication content. For the development of a vehicle the AUTOSAR DEXT is used; for the offboard world typically the ASAM ODX format is used, because it offers higher flexibility across different carlines. Today it is often already a challenge to keep the two worlds consistent. But with the dynamic deployment (offered by Adaptive Platform) it is even more challenging because in worst cases each vehicle has an individual setup of installed SoftwareClusters.
<b>Alternatives</b>	None, because both options are requested by the market.
<b>Remarks</b>	No remarks.
<b>Related requirements</b>	None

## 5.15 Faults inside constructor

<b>Date of approval</b>	2021-06-29
<b>Decision</b>	Calling a constructor that may throw exceptions as part of its defined behavior shall result in a compilation error if the compiler toolchain does not support exceptions. The compilation error shall result from a <code>static_assert</code> with the error message "This constructor requires exception support."
<b>Rationale</b>	Unintended calls to constructors that may throw exceptions are detected at compile time. <code>static_assert</code> is the only viable option. Declaring the constructor <code>protected</code> or <code>private</code> is more complicated. Moreover, <code>static_assert</code> supports a customized error message which explicitly states the cause.
<b>Architectural cluster</b>	<ul style="list-style-type: none"> <li>• Communication Management</li> <li>• Adaptive Core</li> <li>• Intrusion Detection System Manager for Adaptive Platform</li> <li>• Persistency</li> <li>• Execution Management</li> <li>• Time Synchronization</li> <li>• Platform Health Management</li> <li>• Diagnostic Management</li> <li>• Log and Trace</li> <li>• Cryptography</li> </ul>
<b>Category</b>	None
<b>Application affected</b>	Yes
<b>Assumptions</b>	There are toolchains targeted by AUTOSAR, which do not support exceptions.
<b>Constraints</b>	No constraints were identified.







<p><b>Alternatives</b></p>	<ul style="list-style-type: none"> <li>• Constructors that may throw exceptions shall not participate in overload resolution when the compiler toolchain does not support exceptions.                     <ul style="list-style-type: none"> <li>– (Pro) Similar solution as for <code>ara::core::Result::ValueOrThrow()</code></li> <li>– (Con) Changes the overload set. Thus, may result in an unintended change to the program flow instead of a compiler error.                             <ul style="list-style-type: none"> <li>* Unintended changes to the program flow may occur due to overloads and due to conversion functions.</li> <li>* Guidelines regarding constructor overloads might help.</li> <li>* Conversion functions cannot be controlled by the AUTOSAR Adaptive Platform.</li> <li>* The problem does not exist for <code>Result::ValueOrThrow()</code> because the function has no overloads that are available with a toolchain without exception support.</li> </ul> </li> <li>– (Con) May result in lots of <code>#ifdef</code> in vendor-supplied headers.</li> </ul> </li> <li>• Constructors that may throw exceptions shall call <code>abort</code> instead of throwing an exception when the compiler toolchain does not support exceptions.                     <ul style="list-style-type: none"> <li>– (Pro) Constructors that may throw may be used even with a toolchain that does not support exceptions if it can be precluded that an exception is thrown.</li> <li>– (Con) May be difficult to support by vendors, unless they make large-scale changes to their C++ standard library if it does not happen to follow the AR-specified style.</li> <li>– (Con) Unintended calls to such constructors are only detected at runtime and only in the case of an error.</li> </ul> </li> <li>• Implementation-specific behavior                     <ul style="list-style-type: none"> <li>– (Con) Violates [RS_AP_00111]</li> </ul> </li> <li>• All public constructors shall be declared as <code>noexcept</code>. Instead of public constructors that may throw, the named constructor idiom shall be used (even if the toolchain supports exceptions).                     <ul style="list-style-type: none"> <li>– (Pro) Unintended calls to constructors that may throw are detected at compile time.</li> <li>– (Con) Unnecessary restriction when a toolchain is used that supports exceptions.</li> </ul> </li> </ul>
<p><b>Remarks</b></p>	<p>No remarks.</p>
<p><b>Related requirements</b></p>	<ul style="list-style-type: none"> <li>• [RS_AP_00152] Faults inside constructor</li> </ul>

## 5.16 Restart granularity

<b>Date of approval</b>	2021-06-29
<b>Decision</b>	<p>Applications can be restarted in the scope of a <code>FunctionGroup</code>. Ideally, the recovery of supervision errors should be handled in the own <code>FunctionGroup</code>. If the recovery cannot be handled within the own <code>FunctionGroup</code>, it has to be escalated within the <code>State Management</code>. There the coordination for the recovery should take place. This could typically be:</p> <ul style="list-style-type: none"> <li>• the shutdown/restart of multiple <code>FunctionGroups</code>,</li> <li>• the start of other <code>FunctionGroups</code> or</li> <li>• the restart of the entire <code>Machine</code>.</li> </ul> <p>The coordination of the restart of the entire <code>Machine</code> has to be coordinated within the <code>State Management</code> of the platform-core <code>Software Cluster</code>.</p>
<b>Rationale</b>	<p><code>Software Clusters</code> are independently deployable units. They could be added later to the same <code>Machine</code> and then should not harm other <code>Software Clusters</code> (freedom from interference between <code>Software Clusters</code>). Recovery shall always be tried within the <code>Software Cluster</code>.</p>
<b>Architectural cluster</b>	<ul style="list-style-type: none"> <li>• Execution Management</li> <li>• Platform Health Management</li> <li>• State Management</li> </ul>
<b>Category</b>	Safety
<b>Application affected</b>	No
<b>Assumptions</b>	<p>The platform-core <code>Software Cluster</code> is the housekeeping initial <code>Software Cluster</code> which <code>Execution Management</code>, <code>Platform Health Management</code>, and <code>State Management</code> are a mandatory part of (if it is a safety relevant <code>Machine</code>).</p>
<b>Constraints</b>	No constraints were identified.
<b>Alternatives</b>	<p>Applications can be restarted in the scope of a <code>Software Cluster</code>. The <code>Software Cluster</code> is for deployment and not visible in runtime, wherefore it cannot be used in this context.</p>
<b>Remarks</b>	No remarks.
<b>Related requirements</b>	None

## 5.17 Platform-independent development of Software Clusters of category APPLICATION\_LAYER

<b>Date of approval</b>	2021-06-29
<b>Decision</b>	<p>Functional <code>Cluster daemons</code> and their startup coordination shall be part of <code>Software Clusters</code> of category <code>PLATFORM_CORE</code> or <code>PLATFORM</code>.</p>





<b>Rationale</b>	This allows uniform and platform-independent integration of Software Clusters of category APPLICATION_LAYER. Consequently, it shall not be necessary to take care of the platform software when developing an Software Cluster of category APPLICATION_LAYER.
<b>Architectural cluster</b>	<ul style="list-style-type: none"> <li>• Adaptive Core</li> </ul>
<b>Category</b>	None
<b>Application affected</b>	Yes
<b>Assumptions</b>	Market demand is to deliver Machines with pre-installed Adaptive Platform software.
<b>Constraints</b>	No constraints were identified.
<b>Alternatives</b>	Do not make any limitations of platform software. This can lead to a non-uniform integration of the platform software.
<b>Remarks</b>	No remarks.
<b>Related requirements</b>	None

## 5.18 Harmonized logging / tracing

<b>Date of approval</b>	2021-06-29
<b>Decision</b>	Functional Clusters shall standardize their logging/tracing for common use-cases demanded by the market. The standardization shall be for the non-verbose logging/tracing. If applicable it shall be summarized by two standardized tables (one for logging and a second for tracing) listing all standardized log-/trace messages.
<b>Rationale</b>	Standardized logging/tracing within Functional Clusters allows a harmonized evaluation of logging/tracing on vehicle-level.
<b>Architectural cluster</b>	<ul style="list-style-type: none"> <li>• Communication Management</li> <li>• Adaptive Core</li> <li>• Intrusion Detection System Manager for Adaptive Platform</li> <li>• Persistency</li> <li>• Execution Management</li> <li>• Time Synchronization</li> <li>• Platform Health Management</li> <li>• Identity and Access Management</li> <li>• Diagnostic Management</li> <li>• Log and Trace</li> <li>• Cryptography</li> </ul>



△

	<ul style="list-style-type: none"> <li>• State Management</li> <li>• Update and Configuration Management</li> <li>• Network Management</li> </ul>
<b>Category</b>	None
<b>Application affected</b>	Yes
<b>Assumptions</b>	Logging/tracing is necessary for a variety of use cases (root cause analysis, auditing, debugging). Especially, in a distributed environment a harmonization is necessary to enable automated analysis.
<b>Constraints</b>	No constraints were identified.
<b>Alternatives</b>	Do not standardize logging at all.
<b>Remarks</b>	No remarks.
<b>Related requirements</b>	None

## References

- [1] Glossary  
AUTOSAR\_TR\_Glossary
- [2] Explanation of Adaptive Platform Software Architecture  
AUTOSAR\_EXP\_SWArchitecture
- [3] Dynamic Memory Allocation and Fragmentation  
[https://www.researchgate.net/publication/295010953\\_Dynamic\\_Memory\\_Allocation\\_and\\_Fragmentation](https://www.researchgate.net/publication/295010953_Dynamic_Memory_Allocation_and_Fragmentation)
- [4] Dynamic Memory Allocation on Real-Time Linux  
<https://static.lwn.net/images/conf/rtlws-2011/proc/Jianping.pdf>
- [5] TLSF: a new dynamic memory allocator for real-time systems  
<https://doi.org/10.1109/EMRTS.2004.1311009>
- [6] The Memory Fragmentation Problem: Solved?  
<https://doi.org/10.1145/286860.286864>
- [7] C++ Core Guidelines  
<https://github.com/isocpp/CppCoreGuidelines/blob/master/CppCoreGuidelines.md>
- [8] Specification of Adaptive Platform Core  
AUTOSAR\_SWS\_AdaptivePlatformCore