



IIP-Ecosphere Whitepaper

IIP-ECOSPHERE PLATFORM HANDBOOK

Version 0.50

Holger Eichelberger, Amir Shayan Ahmadian,
Andreas Dewes, Marco Ehl, Ahmad Alamoush,
Monika Staciwa, Miguel Gómez Casado

White Paper IIP-2021/005

Disclaimer

The contents of this document has been prepared with great carefulness. Although the information has been prepared with the greatest possible care, there is no claim to factual correctness, completeness and/or timeliness of data; in particular, this publication cannot take into account the specific circumstances of individual cases.

Any use is therefore the reader's own responsibility. Any liability is excluded. This document contains material that is subject to the copyright of individual or multiple IIP-Ecosphere consortium parties. All rights, including reproduction of parts, are held by the authors.

This document reflects only the views of the authors at the time of publication. The Federal Ministry for Economic Affairs and Energy or the responsible project agency are not liable for the use of the information contained herein.

Publication: March, 2023 on <https://www.iip-ecosphere.eu/>

DOI: 10.5281/zenodo.7695124

Gefördert durch:



Bundesministerium
für Wirtschaft
und Energie

Executive Summary

The IIP-Ecosphere platform is a central asset developed by the IIP-Ecosphere project. The core aim of the platform is to research and demonstrate novel platform concepts for Industry 4.0, e.g., asset administration shells as interfaces for software components and resources, unified edge deployment, an AI toolkit or seamless configuration of a platform from network settings via services up to applications running on the platform. This platform handbook provides insights into the rationales, ideas and concepts that make up the design and the realization of the IIP-Ecosphere platform, ranging from an overall layered architecture over a detailed discussion of the design and realization state of each layer up to cross-cutting mechanisms such as the configuration model or the related code/artifact generation.

This platform handbook addresses the technical side of the platform work in IIP-Ecosphere and builds on the intensive prior work on requirements (usage view and functional/quality requirements of the platform). This handbook shall provide means for deeper technical discussions with partners, stakeholders and interested parties, but also allow for a technical understanding to contribute to the platform, e.g., in terms of protocols, platform connectors, services or demonstration applications.

This version of the handbook focuses on the platform release as of March 2023 (version 0.5.0) and supersedes older versions of this handbook/the platform.

Acknowledgements: We are grateful to Dr. Christian Sauer and Alexander Weber from the Software Systems Engineering Group of the University of Hildesheim for cross-reading this document and providing valuable feedback and ideas for improvement.

Contents

1	Introduction.....	7
1.1	Motivation and Goals	7
1.2	Interaction with other initiatives.....	8
1.3	Structure of the document.....	8
2	Tooling and Basic Technical Decisions.....	10
3	Architecture.....	14
3.1	Overview.....	14
3.1.1	Relation to Reference Architectures	20
3.1.2	Stream (Data) Processing	21
3.1.3	Asset Administration Shells	22
3.1.4	Component Interaction Overview	23
3.1.5	Virtual Character of the Platform	26
3.2	Development Streams	26
3.3	Overall Requirements.....	27
3.4	UML Profiles	29
3.4.1	IIP-Ecosphere Profile	29
3.4.2	UMLSec Profile	34
3.4.3	Security and Privacy Profile	35
3.4.4	IoT Component Security and Privacy Profile	39
3.5	Support Layer	40
3.5.1	Asset Administration Shell Abstraction	40
3.5.2	Network Management Support	44
3.5.3	Lifecycle Support	45
3.5.4	System-level Monitoring Support	45
3.5.5	Identity Support.....	45
3.5.6	Resource Support	46
3.5.7	Semantic Id Resolution Support	46
3.5.8	Task Tracking Support	47
3.5.9	AAS Creation and Usage Pattern	47
3.6	Transport and Connection Layer	48
3.6.1	Requirements	48
3.6.2	Transport Component	49
3.6.3	Connectors Component	60
3.6.4	Requirements Discussion	68
3.7	Services Layer	70
3.7.1	Terminology and Background.....	71

3.7.2	Requirements	72
3.7.3	Service Environments	73
3.7.4	Service Control and Management.....	80
3.7.5	Requirements Discussion	87
3.8	Resources and Monitoring Layer.....	88
3.8.1	ECS runtime	88
3.8.2	Device/Resource Management	96
3.8.3	Monitoring.....	101
3.9	Storage, Security and Data Protection Layer	105
3.9.1	KODEX platform service.....	106
3.9.2	Data lakes / Data bases	107
3.10	Reusable Intelligent Services Layer	109
3.10.1	Specific Requirements	109
3.10.2	Data Processing Function Library	111
3.10.3	RapidMiner RTSA service.....	111
3.10.4	Flower-based Federated Learning.....	112
3.10.5	Service candidates ahead	112
3.10.6	IIP-Ecosphere AI Software Service Concept	113
3.10.7	Requirements Discussion	113
3.11	Configuration Layer	114
3.12	Application Layer	122
3.13	Platform Server(s).....	122
3.14	Management User Interface	125
3.15	Test support.....	126
4	Architectural Constraints.....	128
5	Asset Administration Shells	130
6	Platform Configuration Model	134
6.1	Modeling Patterns	140
6.2	Configuration Model Structure	144
6.3	Support for Standardized Connectors/Protocols	145
6.4	Platform Instantiation Process	146
6.5	Container Instantiation	149
6.6	Example Applications	151
6.7	Creating an Application	154
6.8	Project Structures	155
6.9	Default Build Sequences.....	159
6.10	Service Realization Rules and Considerations.....	160

7	Platform Security and Data Protection	164
7.1	Internal Security and Security/Privacy Analysis	164
7.2	Support of the Concepts of the IoT Component Profile in actual Platforms	166
7.2.1	Using the Profile to annotate a system model with proper mechanisms.....	169
7.2.2	Towards an automated analysis to verify required security levels.....	169
7.3	External Security.....	173
8	Implementation.....	175
8.1	Implementation decisions.....	175
8.2	Obtaining the IIP-Ecosphere platform.....	178
8.3	Compiling the IIP-Ecosphere platform	183
8.4	Installing and using the IIP-Ecosphere platform	187
8.5	Pre-build Docker container images	191
8.6	Considerations for a Distributed Server Installation	192
8.7	Environment for Testing and Evaluating the Platform/Applications	193
9	How to apply, extend or contribute	197
9.1	Defining an own application-specific service	197
9.2	Defining an AAS for a device	197
9.3	Implementing a monitoring/alert data service	198
9.4	Extending the platform by adding a component or a platform service.....	198
9.5	Defining a new type in the configuration model.....	199
9.6	Using a different transport protocol	201
9.7	Observe or debug the data processing	201
9.8	Frequently Asked Questions (FAQ)	202
9.8.1	Error parsing HTTP header	202
9.8.2	Maven artifact missing	203
9.8.3	XXX has been compiled by a more recent version of the Java Runtime	203
9.8.4	Platform code cannot be setup in Eclipse, e.g., parent POM missing.....	203
9.8.5	Unknown platform coding conventions	203
9.8.6	Maven does not find app dependencies	204
9.8.7	Execution of application fails due to Java CompileError	204
9.8.8	Services do not start due to problems with javax.el.ExpressionFactory.....	204
9.8.9	Service execution through platform fails	205
9.8.10	Why do platform scripts always check for recent dependency snapshots.....	205
10	Summary & Conclusions.....	206
11	References.....	211

1 Introduction

1.1 Motivation and Goals

The digitalization of the industry increases the effectiveness of technical systems and related processes, but also affects the complexity of the realizing (software) systems. Currently, several approaches are developed in the fields of Internet-of-Things (IoT), Industrial Internet-of-Things (IIoT) or „Industrie 4.0“ (I4.0)¹. To support the industrial transformation towards IoT, IIoT and I4.0, several software platforms were developed that provide different capabilities.

The vision of the BMWi-funded² project IIP-Ecosphere is to enable innovations in the area of industrial production based on connected, intelligent and autonomous systems in order to increase productivity, flexibility, robustness and efficiency of IIoT and I4.0. IIP-Ecosphere aims at creating a novel ecosystem for the “next level” of intelligent industrial production, not only for software-based systems, but also for the people involved in this kind of systems, e.g., automation engineers, software developers, AI experts, startups, venture capitalists, etc. On the software side, one core activity in IIP-Ecosphere is to research and to realize a virtual platform that connects factory installations across companies in a vendor-independent manner. In particular, the platform shall provide easy-to-use access to Artificial Intelligence (AI) in secure and flexible manner.

Towards the design of such a platform, we analyzed in [35] 21 IIoT platforms with specific relevance to IIP-Ecosphere and described in [41, 13] the requirements for the IIP-Ecosphere platform from two different perspectives, namely the usage view and the functional/quality requirements view. The next step is to turn the requirements into an architecture and to implement the platform. The resulting platform shall be open, extensible, vendor-neutral, secure, flexible, configurable, self-adaptive and based on relevant standards as well as on existing Open Source components. In particular, we aim at developing a **virtual platform**, i.e., a platform that utilizes existing, already installed solutions by integrating with them, using accessible output and resources, enhancing them with AI and, if desired, feeding back AI-enhanced information into utilized systems. Thus, we do not aim at replacing existing platforms as those mentioned in [35] rather than enhancing them. Moreover, we aim at demonstrating how research results, e.g., from systematic variability management, security or data protection, can lead to platform concepts that are currently rarely used in IIoT/I4.0 platforms. Besides the desirable abilities mentioned above, following the initial decisions made in [41, 13], the platform shall be service-based and virtualized through containers. One relevant I4.0 standard to integrate the parts and pieces of the platform is the **Asset Administration Shell (AAS)** [32] that we aim to apply as self-description and interface to software components across all platform layers. The consortium discussions regarding a vision of the IIP-Ecosphere platform also emphasized the need to directly communicate with production machines, in particular, to utilize edge devices and, if feasible, cloud technology (such as the upcoming Gaia-X³). This re-shaped the character of the envisioned platform from a purely virtual to a mixed-virtual platform with stronger aspects of a usual IIoT/I4.0 platform, in particular providing **uniform deployment of services** to heterogeneous execution resources such as edge devices, on-premise servers or clouds.

In this whitepaper, we aim at discussing and documenting the architecture and the implementation of the IIP-Ecosphere platform. This happens in an incremental⁴ fashion as, we intentionally mix requirements, architecture and implementation activities in an agile manner. With this approach, we aim at synchronizing the requirements with the architecture and ensuring that the underlying

¹ Translates to some degree to IIoT in German-speaking areas in Europe, partly based on own standards.

² <https://www.bmw.de/Redaktion/DE/Publikationen/Technologie/ki-innovationswettbewerb.html>

³ <https://data-infrastructure.eu>

⁴ Along the realization state, i.e. the releases of the platform software. The version number of this white paper reflects the software release version. Thus, at the beginning some sections may be rather empty.

implementation realizes and fits the architecture. Thus, this document documents the current state at hands, while we aim at updating this document as part of upcoming releases of the IIP-Ecosphere platform. In other words, in this document, we document and discuss the current state of the platform on a feasible level of detail, the underlying implementation, decisions we made and the tradeoffs that we faced. However, depending on the state of the implementation, this document is not meant to be complete but rather to be a “living document” that is updated incrementally. This version of the handbook focuses on the platform release as of March 2023 (version 0.5.0) and supersedes older versions of this handbook/the platform.

This release of the IIP-Ecosphere platform adds the technical prerequisites for the integration of federated learning, automatic creation of platform/application containers, a distributed test/evaluation environment for the platform and applications running on the platform, the automated translation and integration of OPC UA companion specs into the configuration model, LXC as alternative container manager, an integrated build and installation process for applications and the platform (as inspired by a service development workshop), an upgrade to BaSys 1.3.0 and many more detailed features and improvements⁵. Further, an initial video tutorial on service development for the platform based on a project-internal workshop have been created and published⁶.

It is important to mention that this document is also meant to be a basis for discussions with the respective teams in IIP-Ecosphere⁷ (mainly Think Tank “Platforms” and KI-Accelerator) and with all kinds of platform stakeholders in order to help, improve, influence or integrate with development of the IIP-Ecosphere platform. So far, it helped to onboard various co-workers and stimulated detail decisions and clarifications.

1.2 Interaction with other initiatives

Work on the IIP-Ecosphere platform is influenced by interaction with other initiatives, in particular

- The IIP-Ecosphere IIoT platform overview [35] indicating challenges and potential for future AI-based I4.0 platforms.
- Interactions with other funded projects: DaPro⁸, BaSys⁹, FabOs¹⁰, Service-Meister¹¹.
- Internal IIP-Ecosphere stakeholders with interest in validating the platform (in conjunction with their own approaches) or for contributing components, e.g., dedicated services. In the remainder of the funded time of IIP-Ecosphere, the “AI accelerator” work package plans to make contributions in terms of customizable (AI) services. First steps in this direction have been done, e.g., in terms of feasibility studies with the IIP-Ecosphere demonstrators or a collection of candidate services for generalization.

1.3 Structure of the document

A typical first section of a platform handbook could be a summary of the requirements to be realized. As stated in Section 1, the IIP-Ecosphere team summarized the results of the requirements collection for the platform in two other whitepapers, namely the usage view [41] and the functional/quality requirements view [13]. For pragmatic reasons, these two documents have been prepared partially before and partially while designing the platform architecture, so that they are synchronized with the

⁵ <https://github.com/iip-ecosphere/platform/blob/main/platform/documentation/RELEASES.md>

⁶ <https://www.iip-ecosphere.de/angebote/plattform/> and https://www.youtube.com/playlist?list=PL5VSYtiD_PfdxUDXGIX53UsHdQIXcHYK7

⁷ See <https://www.iip-ecosphere.eu/> for a summary of the project structure.

⁸ <http://dapro-projekt.de/>

⁹ <https://www.basys40.de/>

¹⁰ <https://www.fab-os.org/>

¹¹ <https://www.servicemeister.org/>

work described here. In order to avoid inconsistencies, we are not repeating the requirements in this document rather than referring to [41, 13] through requirements identifiers defined there.

In Section 2 we introduce the tooling that is used for developing the architecture model and the implementation. A brief discussion of the tooling and the rationales for certain decisions is relevant at that point as the decisions significantly interact with the modeling concepts, i.e., affect the set of concepts that we practically can use for specifying, describing or realizing the architecture. Moreover, Section 2 already indicates that the ultimate outcome of our work is not “just” an architecture rather than an implemented and working platform.

In Section 3 we introduce and discuss the architecture of the IIP-Ecosphere platform, ranging from the UML profiles used, over the lower transport up to user-defined applications. This section is not only intended to present the architecture as it was designed rather than also the tradeoffs that we faced and the decisions that we made towards the actual architecture. In Section 4, we summarize architectural constraints that must be obeyed by the implementation. In Section 5, we discuss the representation of the platform components in terms of Asset Administration Shells, which are used as a uniform way to represent interfaces and communication among components.

One aim of the platform work in IIP-Ecosphere is to research concepts on systematically and consistently configuring such a platform, ranging from network settings over available resources or services up to the wiring of re-usable parts and pieces to IIoT-applications. In Section 6, we elaborate the structure of and the concepts of the model to specify decisions that must be made to turn alternative or generic components into an installable platform with user-defined applications. We will also discuss, how to utilize such a model, not only to validate configuration decisions, but, in particular, to automatically generate platform instances, artifacts or glue code as one means of supporting platform users to create IIoT-applications.

In Section 7, we discuss mechanisms ensuring the security of the platform. In Section 8, we detail how to obtain, install, instantiate and use (depending on the implementation state) the IIP-Ecosphere platform. In Section 9 we summarize steps on how to extend, contribute to or use the IIP-Ecosphere platform.

In particular for Sections 3 to 9 it is important to recall that the IIP-Ecosphere platform is currently under agile and incremental development, i.e., while some sections are detailed and an implementation is provided for the respective components, other components are still in planning and not yet realized. The design and implementation state will change and evolve over time as the architecture and the implementation will do. To detail the respective realization state, we will refer to the requirements in terms of realized, modified or, if needed, even deferred or excluded requirements.

Ultimately, in Section 10 we will summarize and conclude this document. In Section 11 we list references to other work that we rely on.

2 Tooling and Basic Technical Decisions

Tooling is an important topic when creating an architecture and when implementing it in terms of executable code. In this section, we briefly describe the tooling decisions made by the involved partners, as they affect the available options for modeling the architecture and for realizing it.

The **architecture** is designed using the Unified Modeling Language (UML) [30]. We will not introduce UML in this document rather than assuming that the reader is sufficiently familiar with UML. As tool support, we use Eclipse Papyrus¹². While there is a broad range of modeling tools available, in particular commercial ones, we decided to use Papyrus for two major reasons:

- 1) During architecture modeling concepts for security and data privacy shall already be integrated and the architecture shall be evaluated in this direction. Therefore, we will use UMLsec [24] as well as a specific security profile developed for IIP-Ecosphere. UMLsec has been successfully applied with Papyrus and with the Eclipse UML modeling tools, advocating Papyrus/Eclipse as a natural choice for our work. For applying the security concepts, the respective UML profiles must be installed and integrated into the model. For an automated security analysis, the additional Eclipse-based CARISMA¹³ tool must be installed.
- 2) In contrast to commercial software, Papyrus is available to the interested public as it is released as Open Source. This facilitates platform work, as we plan to release the UML model of the IIP-Ecosphere platform as part of one of the platform releases. Moreover, as it is based on Eclipse, further available tools and model translations from the Eclipse ecosystem may be utilized.

Although Papyrus offers various UML modeling capabilities, in particular the behavioral modeling for state machines, sequence or communication diagrams are currently not completely stable. This, however, affects the available options and concepts for modeling the platform architecture. Thus, in some cases, more recent modeling concepts could have been used that are not available for this reason. Unfortunately, the realization state of Papyrus also affects the layout of the included diagrams, e.g., if technical screen resolutions change, which could be presented in more pleasing manner if only some more diagramming functionality would be available. This is also true for the Papyrus diagram export, which so far produces only formats (bitmap, SVG) that unfortunately can only hardly (or with some inconvenient transformation steps) be used with Microsoft Word. Thus, we include UML figures taken from the architecture model as bitmaps into this document.

Along with the architecture and the design of individual components, also architectural constraints arise, e.g., that in particular for alternative components, dependencies to underlying libraries must be private to the respective platform component, i.e., and globally used by other platform components. We will discuss the architectural constraints of the IIP-Ecosphere platform in Section 4 as a specific summary of the architecture section. Section 3 may already indicate or mention such constraints.

For **implementing** the architecture, we must integrate existing components and consider that in particular AI services will be realized in different programming languages.

- For the **Java** components, we rely on Eclipse (so far 2021-03, version 4.19.0) with Maven¹⁴ (version 3.6.3), Git¹⁵ and checkstyle¹⁶ (version 8.40) integrations. Fundamental technical decisions are documented along with the code. As we use Maven for the platform installation,

¹² <https://www.eclipse.org/papyrus/> version 4.8

¹³ <https://rgse.uni-koblenz.de/carisma/>

¹⁴ <https://maven.apache.org/> also tested with 3.8.5

¹⁵ <https://git-scm.com/>

¹⁶ <https://checkstyle.sourceforge.io/> known working versions are 8.35 up to 8.41

a Java Development Kit (JDK) is required rather than a Java Runtime Environment (JRE). We just mention some of the decisions here: The dependency management and the build process are specified in Maven. Templates for code formatting and validation of the formatting are available for checkstyle in the source code repository as part of the Eclipse project for managed platform dependencies and shall be applied prior to any commit. A common logging framework was selected (slf4j) based on decisions of components to be integrated. Components of the IIP-Ecosphere platform are represented as individual Eclipse projects. For compliance with yet unknown edge devices, we required until version 0.4 that (at least the lower, edge-related) layers are executable with Java 1.8 (as this is also the case for many available IoT libraries). Due to recent decisions of the AAS reference implementation BaSyx to provide binary components for Java 11 (as well as test broker implementations also switching to or requiring Java 11), we decided to abandon the original Java 1.8 requirement in version 0.5 of the IIP-Ecosphere platform and to rely on Java 11 as minimum/build Java version. For the continuous integration, the build/deployment process is specified due to technical reasons in ANT and for platform examples partially as Bash shell scripts.

- While some AI methods may also be realized in Java, nowadays AI methods are frequently implemented based on **Python**. For Python services (as for Java-based services), a service execution environment is provided, which is responsible for the communication with related Java components, so that an AI developer does not have to work with both languages, protocol details or a plethora of different protocols. For the service environment, we rely on Python 3.9.6¹⁷, a rather recent version as modern AI frameworks often also require a recent Python version. Python services musts explicitly declare their dependencies, e.g., used AI frameworks and state this information in the platform configuration for automated creation of installation artifacts, in particular containers.
- In particular, for the initial versions of the platform, we prioritize **dependency reduction and binary image size** over alternative, potentially more modern programming approaches. With this decision, we also try to reduce the risk of incompatible libraries for platform components that still need to be integrated. Thus, we decided not to use a basic framework like Spring as foundation as it may leads to later conflicts (as we experienced later for Spring Cloud Stream and the AAS reference implementation Eclipse BaSyx). Akin, we prefer in some places boilerplate code over annotation-style programming, e.g., in platform parts where a later revision with yet unclear external decisions can be foreseen (standardization of AAS, future development of BaSyx).
- Some components require **technical settings** for their startup, e.g., certain internet addresses or basic security certificates to announce the own instance, to request or contribute information. The aim is to reduce such explicit setup information to a minimum as it is a source for inconsistencies. For this purpose, such information shall be managed centrally, instantiated into binary components or distributed via discovery protocols where feasible. So far, as alternative, no automated discovery mechanisms (for I4.0) settings was suggested/integrated, which could ease the setup. Further information not required to startup a component shall be made available via the (joint) AAS of the platform. Technical settings that may be subject to modifications by administrators shall be represented in a uniform and human readable manner. For stored setup information we rely on Yaml¹⁸, for machine-readable complex data in AAS on JSON¹⁹. Regarding terminology, we distinguish between **Setup** (the technical information, e.g., in Yaml, in practice often also called configuration) and the **Configuration**

¹⁷ Python 3.8 does not support the automatic class loading employed by the platform.

¹⁸ <https://en.wikipedia.org/wiki/YAML>

¹⁹ <https://www.json.org/json-en.html>

(the managing part, potentially generating the setup information for consistency). Related source code shall be named accordingly²⁰.

- Components shall internally communicate via **interfaces** in order to reduce (accidental) dependencies. Alternative and optional components shall be realized as a kind of plugin and register themselves into the platform. On the Java side, we rely on the Java Service Loader (JSL) mechanism, which associates concrete implementations to their respective (descriptor) interfaces. The relation happens through a specific form of file that is evaluated by the JSL mechanism upon request. We use that mechanism to define, e.g., factory instances, to compose AAS but also to set up the component lifecycle, e.g., to handle the start and shutdown process.
- So far, no mechanisms to shield (the dependencies) of individual platforms against each other was necessary, as, e.g., technical dependency conflicts could be successfully resolved through global version restrictions. However, we are aware of the fact that in particular through external contributions, conflicts may arise that cannot be solved in this manner. Thus, for future releases, we plan to investigate, whether approaches like **OSGi** (Open Services Gateway Initiative) could help to avoid unintended or unexpected conflicts.
- All components shall provide sufficient **tests** for their functionality. Tests shall be executed during the **continuous integration** (CI) of the platform and also usual test metrics shall be recorded. **Test artifacts**, e.g., setup files created specifically for testing components or dependencies used only for testing, must be **strictly separated from production code**, e.g., reside only in test resource folders. In particular for Java components this is important as setup files that are accidentally placed in production resource folders may take precedence on the classpath over generated setup folders, i.e., prevent that the configuration decisions made by the user are enacted.

As stated in Section 1, for several reasons one objective of the IIP-Ecosphere platform is to use existing Open Source solutions wherever feasible. However, not all **Open Source licenses** are per se permissible in industrial contexts. Therefore, the IIP-Ecosphere consortium has reviewed Open Source licenses and categorized them into four categories:

- 1) Usable without limitations, e.g., MIT, BSD-2-Clause, BSD-3-Clause, ISC, CDDL1.0, Eclipse-Dist-1.0.
- 2) Permissible, but potentially problematic, e.g., Apache 2.0, LGPL-2.1, Artistic-1.0-Perl, EPL-2.0, MS-PL, MPL-1.1.
- 3) Commercial licenses.
- 4) Not allowed or at least problematic, in particular due to copy-left implications, e.g., GPL-2.0, GPL-3.0, EPL-1.0, AGPL-3.0. In some cases, the use of binary artifacts of software under such licenses may still be permitted as long as license information and the origin are stated and the underlying code is not modified or included.

These categories shall be considered already during the design of the IIP-Ecosphere platform and may effectively limit potential candidates. Licenses of the first two categories may be used (with care), the remaining shall be avoided. This is in particular true for platform components that constitute mandatory core functionalities of the platform. Commercial licenses may be used depending on the decision of the installing organization. Components relying on commercial licenses shall be optional by

²⁰ Initially, we aimed for an alignment with Spring, also calling the technical setups a “configuration”. However, this led to some confusion, so we decided for version 0.3.0 to refactor the platform code according to the setup/configuration naming convention introduced above. However, some parameter/variable names and comments may still use configuration, config or cfg where setup would now be correct. We will try to clean up these (local) inconsistencies incrementally over the time.

default and, thus, their use is the decision of the using organization. Analogously, also software under not permissible licenses could be used in optional parts of the platform, but to avoid later license conflicts, licenses of category 4 shall be avoided wherever possible.

The source code of the IIP-Ecosphere platform is made publicly available in the **GitHub** space of IIP-Ecosphere²¹. Moreover, to foster transparency, the development of the IIP-Ecosphere platform happens in public. In later stages also the underlying architecture model shall be made available to support external and future developments after the project lifetime. As far as possible, components are subject to CI using the Jenkins server of the Software Systems Engineering (SSE) group at the University of Hildesheim. Upon successful builds, artifact snapshots are deployed by the CI processes to the Maven repository²² of the SSE group. Java parts of stable releases are deployed to Maven central²³.

²¹ <https://github.com/iip-ecosphere/platform/>

²² <https://projects.sse.uni-hildesheim.de/qm/maven/>

²³ E.g., <https://repo1.maven.org/maven2/de/iip-ecosphere/platform/>,
<https://search.maven.org/artifact/de.iip-ecosphere.platform/transport>

3 Architecture

The architecture of the IIP-Ecosphere platform aims at realizing the requirements collected in the project [41, 13] in terms of software. In this section, we discuss the design of the individual parts and components of the platform. Please note that as mentioned in Section 1, we follow a pragmatic agile approach to the development of the platform, which involves forward and feedback cycles among requirements, architecture and implementation. Thus, depending on the realization state, not all platform components may be completely described in this version of the document, i.e., we will work out sections incrementally depending on the realization state.

We start in Section 3.1 with an overview of the platform layers and dive then into their details. At the end of Section 3.1, we detail some further basic aspects, namely relation to reference architectures in Section 3.1.1, the concept of data flow processing in Section 3.1.2, a brief introduction into asset administration shells in Section 3.1.3, high-level component interactions in Section 3.1.4, and the virtual character of the platform in Section 3.1.5. Section 3.2 indicates the coarse-grained development streams. Section 3.3 takes up the general requirements from [13] as context for the platform architecture. As basis for the architecture description, we discuss in Section 3.4 the used UML profiles and go through the layers of the infrastructure, first as overview and then one section per layer, starting at the bottommost layer.

3.1 Overview

The overall architecture of the IIP-Ecosphere platform follows a layered style (see Figure 2 with only high-level relations shown) based on components and services (R4 in [13]). As far as feasible, we aim for a strict (logical) layering, so that for two adjacent layers l_i and l_u (with as “the lower layer” l_i being located below “the upper layer” l_u), only l_u (and not its transitive upper layers) shall access or call l_i directly. Moreover, there are also aspects that cross-cut visibly or invisibly in this layered structure.

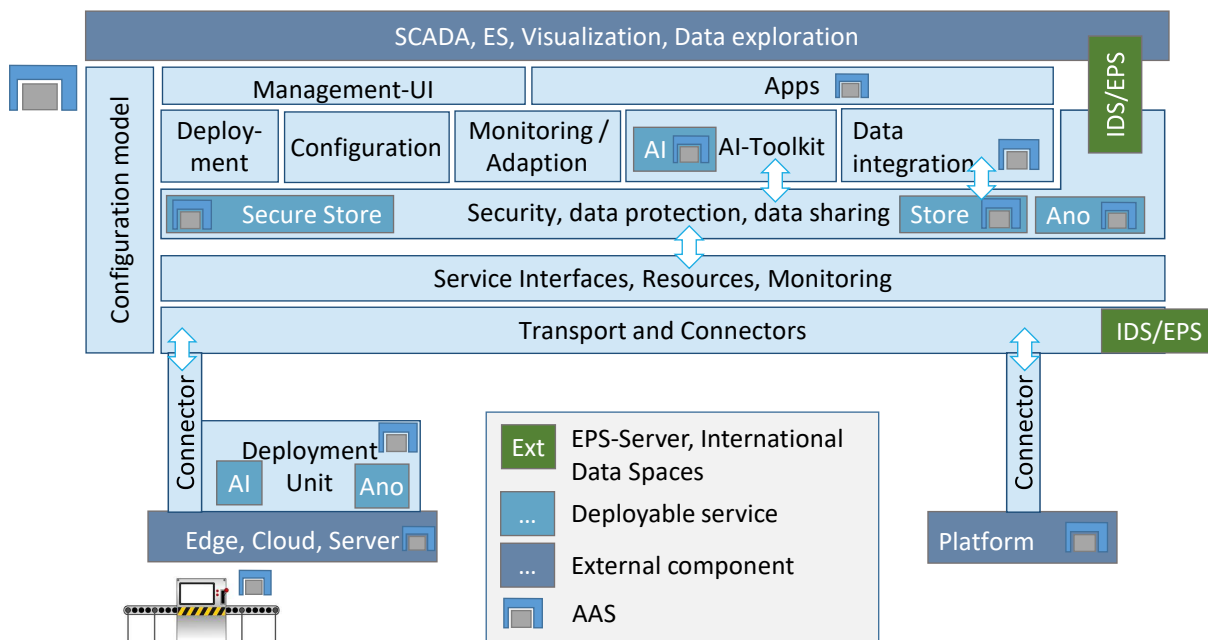


Figure 1: Main platform components as block diagram.

Figure 1 depicts an overview of the high-level building blocks of the IIP-Ecosphere platform. The overall goal is to enable service-based AI-enabled applications to be deployed to available resources utilizing standardized protocols. Therefore, the platform encompasses on the two lower layers (middle of Figure 1) the platform-internal data transport, connectors to external data, e.g., to machines, basic service interfaces, resource abstraction and monitoring facilities. These two layers form the basis for

units that can be deployed to edges, servers or to a cloud. The remaining layers add services, e.g., on data protection, data integration, AI, runtime adaptation/centralized monitoring, overall configuration and deployment of services/containers. If permitted, services from those layers can be deployed to available resources by the platform. In contrast, Figure 2 focuses on the software side and the actual layering.

- **Asset Administration Shells** (R7, [32]) are used in the IIP-Ecosphere platform in two forms: On the one side, AAS represent assets created by different vendors (e.g., a machine, an edge device, an already installed platform, a certain service, a storage mechanism or an App composed of services). The respective AAS may be provided by a third party. Also, the platform itself forms such an asset that deserves an own AAS. On the other side, we utilize the mechanisms of AAS also for describing interfaces of individual interacting components within the platform, but also for realizing this interaction. These components may be internal or external, i.e., also interfaces provided by external AAS may be used. In particular, AAS (submodels) for internal components may be created for the purpose of internal communication rather than external component realization and, thus, may not follow official standardized formats²⁴. For the realization, we rely on the “reference implementation” BaSyx²⁵. An integration of AAS as well as support for realizing (internal) Asset Administration Shells in IIP-Ecosphere style will form the bottom-most layer of the platform.
- The platform contains an event-based **transport messaging** mechanism, e.g., a Broker, so that components and services can communicate among each other independent of the layering. Although this implies certain degrees of freedom and may be used to bypass R7 in exceptional cases, the event-based messaging shall not happen in an ad-hoc or chaotic manner undermining the layer structure. Further, uncontrolled messaging may accidentally overload the broker(s), in particular if the broker is involved in the processing of soft-realtime data streams (one potential manifestation of R10 [13]). As event-based communication and data streaming are essential to the platform, they occur on one of the fundamental layers (Transport) utilizing the external (abstract) components Broker and StreamingLibrary.
- **Variability management** and **consistent configuration** typically do also cross-cut layers, as variability instantiations may affect all components. This is already reflected in the requirements, where *configuration model* occurs in many different functional topics, see e.g., also for implicit information R8, R19f, R20, R28, R30, R31, R34, R40-R43, R62, R64, R73, R77, R80, R86, R89, R93-R101, R104, R107, R112, R119-R122, R131, R134 in [13], but also in the (variability-based) configuration model that crosses several building blocks/components in Figure 1 and Figure 2. Moreover, some layers require access to the configuration, in particular at runtime, e.g., to determine whether migrations of components are needed or how adaptations shall be enacted. However, also here a chaotic use of the configuration can easily lead to unmanageable dependencies. Therefore, we modularize the configuration along the layers (as indicated in Figure 2), and, if required, provide access to the individual configuration modules. Similarly, only some few selected mechanisms to instantiate variability shall be utilized, in particular code generation, generation of setup files and artifact selection while packaging.

For short, the layers of the platform from bottom to top:

²⁴ At the point of writing, several forms of AAS are in standardization, but most known to us do not aim at platform components. Wherever possible, we utilize existing standards, e.g., for device nameplates, or try to adopt the style of related standards to express proto-AAS, e.g., for software services.

²⁵ <https://www.eclipse.org/basyx/>

- Support Layer:** The support layer (not shown in Figure 1) realizes basic abstractions and helpful functions for the upper layers of the platform. The core aim is to reduce repetitions of non-trivial management functions or functions to create common AAS structures and to foster internal conventions, e.g., how to represent certain information in AAS. Moreover, it contains an abstraction of the underlying AAS implementation, serving for both, more flexibility (allowing to also use other implementations) and risk reduction.

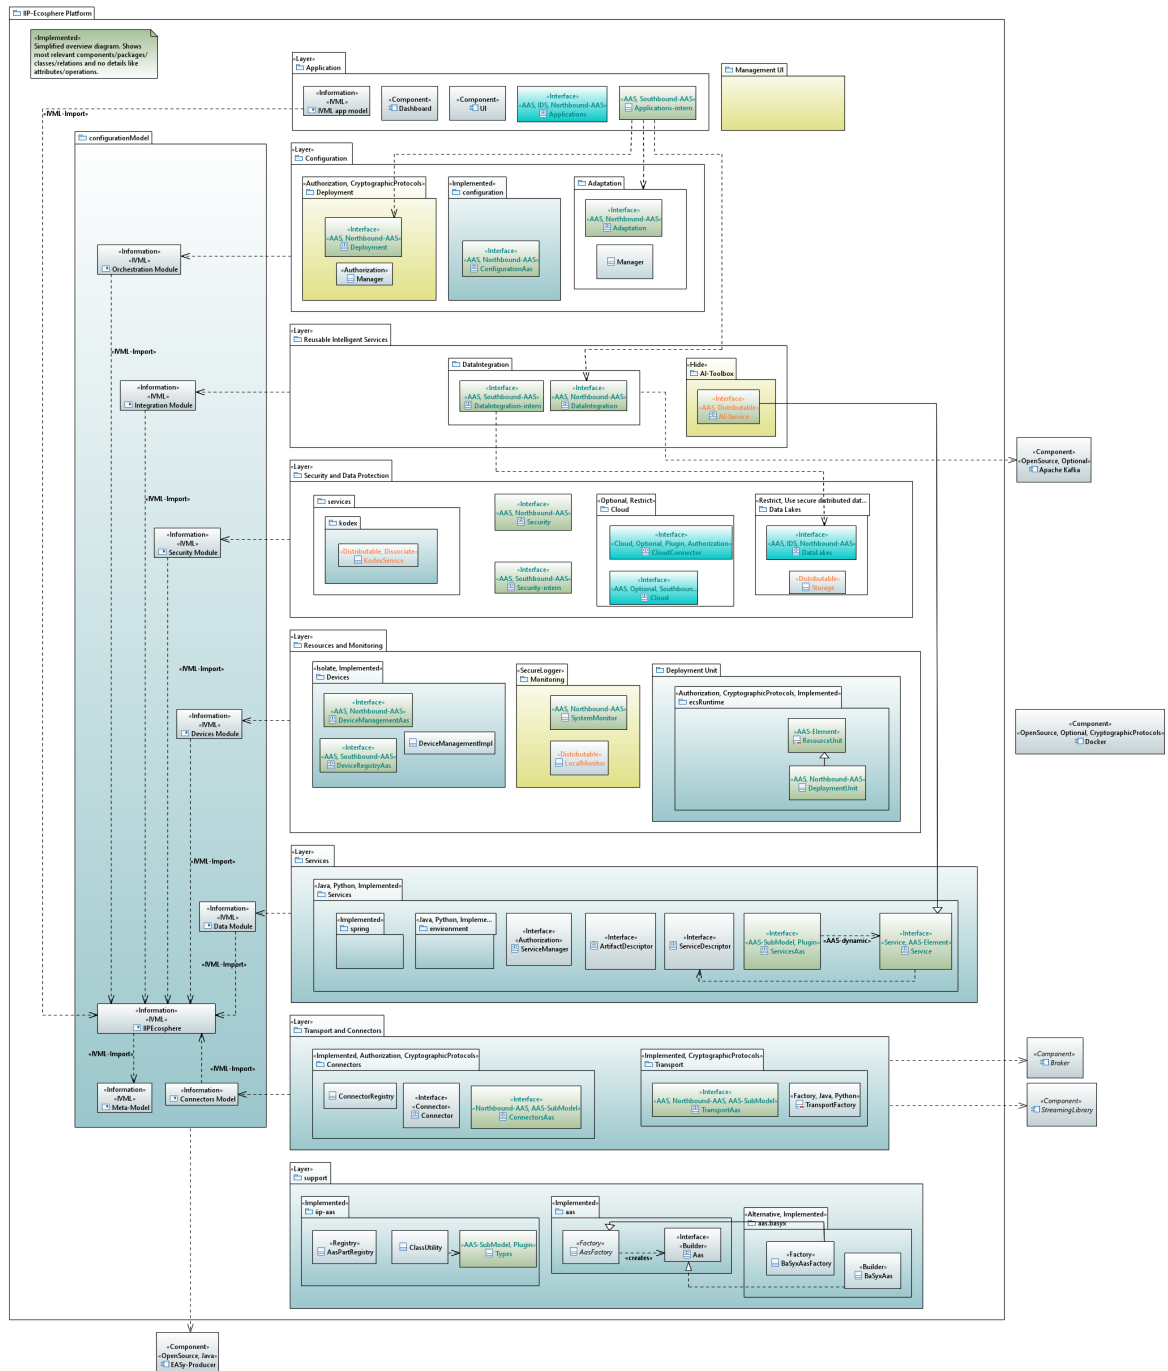


Figure 2: Layered platform overview with indicating only relevant high-level relationships²⁶.

²⁶ Colors indicate the realization state and element categories. Green components indicate AAS components, turquoise layers/components are actually realized (at least in an initial version), red parts are so far not realized

- Transport and Connectors Layer:** This layer is responsible for connecting devices among each other and with platform services using appropriate protocols and formats from the I4.0 domain. However, several protocols and formats impose different tradeoffs in functionality, performance, security and legal/normative impact. This layer integrates such protocols in a flexible manner. The role of the **Transport Component** is to abstract over relevant protocols such as MQTT²⁷, AMQP²⁸, or OPC UA pub/sub²⁹ to integrate the abstraction with the technology used for streaming (StreamingLibrary) and to provide an environment for protocol/format connectors. In contrast to recent platforms [35], where a single fixed transport protocol is not uncommon, we want to avoid making such basic decisions on behalf of the user already on this layer. Further, for the streaming technology several candidate approaches with their tradeoffs are known. The idea is to prepare a flexible integration and to link this decision to the selected transport protocol. Similarly, connections to production machines and already installed platforms are abstracted by the **Connectors Component**. Such a Connector may utilize similar protocols as the Transport Component, but also protocols at higher semantic levels such as OPC UA providing an own information model shall be made available. In contrast to the Transport Component, which passes through given data, here only subsets of the data being available to a connector may be ingested in the platform and information/commands originating from the platform may be transported back, e.g., to reconfigure an underlying machine. The Connectors may optionally include functionality of the International Data Spaces (IDS)³⁰ for secure access to data.
- Services Layer:** Openness and extensibility through services of different kinds, in particular AI services, are at the heart of the IIP-Ecosphere platform. To be useful for an application, services must be parameterized and orchestrated, e.g., their data (streams) must be connected to other services or connectors. While the interconnections will be handled by the Transport and Connectors Layer, the Services Layer defines the basic service interfaces (Services) as well as the services execution environments, e.g., for Java and Python. Services may be realized in different programming languages and, thus, demand different integration capabilities, ranging from direct calls (Java services) to operating system processes (Python services in an IIP-Ecosphere execution environment, GO, or even standalone Java programs). Therefore, this layer also encompasses service environments for the specific languages or basic, re-usable integration components.
- Resources and Monitoring Layer:** To become effective, services must be deployed to resources/devices (in terms of a Deployment Unit) and monitored at runtime. In IIP-Ecosphere, deployment targets such as edge devices shall describe themselves in terms of AAS and perform a registration with the device registry (Devices), which reflects its data into the runtime structures of the platform. For deployment, the Deployment Unit (more precisely, the ECS runtime from [41]) receives commands via its AAS from the platform, downloads a container including the service implementations³¹ and starts the container. Also the execution of the services in the container must be monitored, which may involve reusable monitoring probes provided by the platform as well as application-specific probes. The reusable mechanisms are provided by the Monitoring component, which (in terms of probes and

and may finally even be omitted (e.g., functions of semantic mapping and routing are already taken over by other components) and orange parts are currently in realization.

²⁷ <https://mqtt.org/>

²⁸ <https://www.amqp.org/>

²⁹ <https://opcfoundation.org/news/press-releases/opc-foundation-announces-opc-ua-pubsub-release-important-extension-opc-ua-communication-platform/>

³⁰ <https://www.internationaldataspaces.org/>

³¹ Assembling the containers is managed by the Configuration Layer as described below.

signaling) is part of the service environment while the aggregation of the monitoring data happens on central IT level. The Monitoring component also uses the capabilities of the support layer (monitoring in terms of AAS) and the Transport and Connection layer (fast track signaling, alarms) and may issue alerts in generic as well as application-specific manner to further layers.

- Storage, Security and Data Protection Layer:** Security and data protection in the IIP-Ecosphere platform encompass of two parts, 1) cross-cutting mechanisms that can be used to implement security and data protection in any component, e.g., authentication, and 2) centralized or distributable mechanisms to support security and data protection, e.g., services supporting data protection or data storage. While the cross-cutting mechanisms occur in all layers (directly or indirectly controlled through the platform configuration), this layer primarily focuses on the second part. Thus, it provides access to the overall security configuration, e.g., authentication tokens or cryptographic keys for accessing edge devices. Further, this layer realizes components (optionally) enhancing the security and data protection, e.g., stream-based services for Anonymization and Pseudonymization, external (Cloud) communication connectors and (optionally secure) Data Lakes. Data lakes/stores/databases may be distributable components to be packed into deployment units, e.g., to buffer data on edge devices.
- Reusable Intelligent Services Layer:** The components described so far (as well as not mentioned administrative services provided by the platforms) can be used to develop simple applications similar to existing platforms [35]. This layer shall pave the way for open, extensible and reusable intelligent services. The Data Integration collects data from running services (as defined during the orchestration) and integrates the data with additional information such as floor plans, order data etc. The integrated information may be stored in storages provided by the Data Lake component(s). The actual functionality of this component in the context of a running application is also defined in the platform configuration. Finally, the AI-Toolbox shall contain re-usable AI services that can be parameterized and orchestrated to form a running application. Basic functionality for this layer is part of this release, e.g., the Python service environment or an optional integration of the RapidMiner Real Time Scoring Agent (RTSA) as generic, re-usable AI service.
- Configuration Layer:** The configuration layer contains components to manage the platform configuration. The Configuration component is responsible for composing reusable and application-specific services and representing the information in terms of the application specific-modules of the platform configuration. The Deployment component is responsible for deciding which services shall be executed by which device (e.g., edge, server or cloud) depending on runtime information available in the platform configuration. Based on these decisions and device-specific information provided by a device AAS, deployment containers are created automatically and made available. Furthermore, the Deployment component shall take the dynamic state of the platform reflected in the platform configuration into account to optimize the container/service deployment at runtime, e.g., supported by generated service glue code or dynamic re-routing of data by the Transport and Connection Layer or the Streaming Library. In addition, the Adaptation component is responsible to decide about configuration changes to deployed services as well as selection of alternative services at runtime (supported by similar mechanisms as for runtime deployment adaptation).
- Applications Layer:** Applications are described by configuration modules and may ship with application-specific components, e.g., AI services or monitoring probes. Although not visible here, glue or transport code generated for orchestrated services implicitly belongs to the applications. The execution of the applications shall be visualized by (as far as feasible) generic

Dashboard components. Further, external AAS-based access to selected aggregated information of the platform can be made available through secure mechanisms, e.g., IDS.

- **Management User Interface:** Ultimately, a (simple) platform Web user management interface (UI)³² relying in particular on components of the Configuration layer, the AAS of the platform as information sources as well as AI-enabled applications run on top of the platform. It is important to emphasize, that although the management interface is realized as a Web UI, the platform must not necessarily be installed/deployed in a Web/Cloud setting, i.e., on-premise installation and use of the Web UI via a browser is one important installation alternative for the IIP-Ecosphere platform.

International Data Spaces (IDS) [22] is a virtual data space leveraging various standards, technologies, and governance models to enable secure and standardized data exchange in a trusted environment. IDS offers a decentralized data storage where several companies share data through IDS Connectors. Moreover, IDS allows to deploy various internal and external applications into the IDS Connectors in order to provide various services on top of data exchange processes. Furthermore, IDS introduces a so-called security profile indicating the capabilities of a Connector to maintain this secure and trusted. As discussed above, security is usually cross-cutting, i.e., while individual mechanisms may enhance or wrap IIP-Ecosphere platform connectors, e.g., to act as IDS connectors, other mechanisms may be more on the central side, such as an integration with the IDS data storage.

Gaia-X³³ aims to form data spaces for companies and citizens to collate and share data in a way that the users can keep control over them. In particular, the users shall decide what happens to their data, where it is stored, and always retain data sovereignty. The architecture of Gaia-X is based on the principle of decentralisation and integrates a multitude of individual platforms that all follow a common standard.

An IIP-Ecosphere installation may interface with both, IDS and Gaia-X, one of them or even none of them as desired by the user, e.g., to remove respective connectors and components completely from the individual platform instance upon platform instantiation, if they are not desired. For access and privacy protection, the IIP-Ecosphere platform may employ also different approaches, e.g., a privacy shield like the KIPROTECT Endpoint System (EPS)³⁴.

The full stack shown in Figure 2 is not required for all kinds of installations. E.g., on a resource such as an edge device, a cloud or a server, a **specialized runtime** is needed (ECS runtime from [41]) to take control over containers and services. The ECS runtime can be composed from a subset of the layers as indicated in Figure 3. The basic layers such as Support as well as Transport and Connectors must be present (from the Support Layer also mechanisms for dynamic network management). For managing containers, at least the deployment unit from the Resources and Monitoring Layer is needed. However, the Services Layer is optional for an ECS runtime, at least in the same container. If an ECS runtime installation also ships with all dependencies needed to run the configured services (e.g., Python and AI libraries), then it might make sense to also have the service manager from the Services Layer present. Otherwise, the Services Layer shall optionally be executable in an own container, based on the Support as well as Transport and Connectors Layers. This container would then be under the control of the ECS runtime, i.e., the local Resources and Monitoring Layer.

³² As discussed in [13], user interface and dashboards are formally out of scope of our funding contract. However, if feasible, we plan to realize at least a simple (Web) user interface in one of the next releases.

³³ <https://www.data-infrastructure.eu/>

³⁴ <https://github.com/kiprotect/eps>

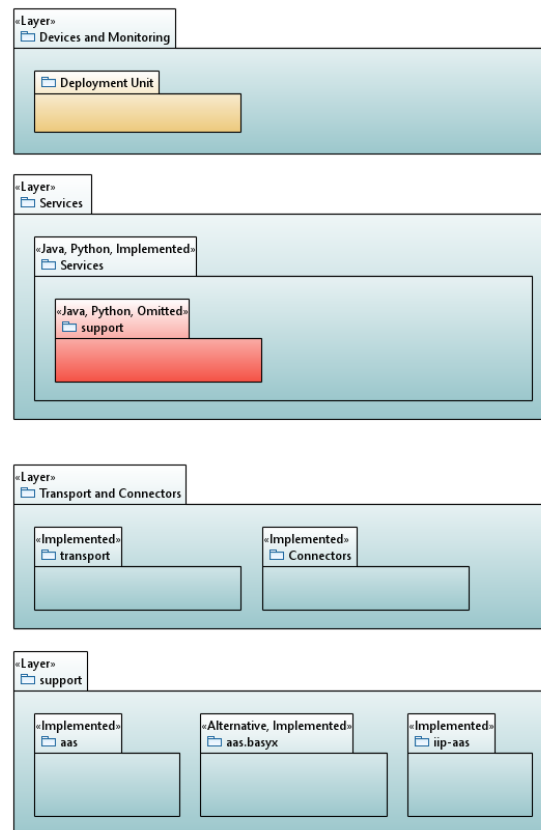


Figure 3: Layers and components required to build an ECS runtime.

The **Java 1.8 restriction** stated in Section 2 applies in particular to the layers shown in Figure 3 in order to enable compliance with unknown edge devices. Although execution on edge devices shall be virtualized in terms of containers (see [13], in particular R30), it may also be required in some settings that the ECS runtime is directly executed by a Java virtual machine on the edge device. As far as we could see at the point in time when designing the architecture of the IIP-Ecosphere platform, this is no significant limitation as relevant (client) libraries for AAS, IoT protocols or connectors or data streaming can be used there. However, we are also aware of the fact that in particular for testing further (broker) libraries may be required, where e.g., the restriction to Java 1.8 may not be fulfilled. This can be mitigated to some degree, as in an installation also equivalent functionality in terms of native programs or other programming languages are available and may be used. As also stated in Section 2, this constraint may be relaxed for the remaining layers shown in Figure 1.

3.1.1 Relation to Reference Architectures

IIP-Ecosphere aims at interrelating and adhering to reference architectures such as RAMI 4.0 [33]. Although we use an own naming of the platform layers, they map nonetheless to layers defined by RAMI 4.0 as summarized in Table 1. However, it is important to recall that the IIP-Ecosphere platform shall be a virtual platform, i.e., it shall in particular (be able to) build on existing installations without implementing a complete IIoT platform. Thus, it is not relevant to meticulously adhere to all RAMI levels, in particular not to the lower levels targeting field devices (as already scoped out in [41, 13]). In addition, our architecture includes some (crosscutting) layers that do not directly fit into the picture of RAMI³⁵, but are important to operations, research and contributions of IIP-Ecosphere.

³⁵ Crosscutting aspects are better covered by IRA [23].

Table 1: Mapping RAMI 4.0 and the IIP-Ecosphere architecture

RAMI 4.0 Axis	RAMI 4.0 Level	IIP-Ecosphere Layer/Component
Layers	Asset	Not in scope [41, 13], represented through edge AAS
	Integration	Support Layer, Transport and Connectors Layer
	Communication	Services Layer
	Information	Reusable Intelligent Services Layer
	Functional	Application Layer
	Business	On top of Application Layer via Applications AAS
Hierarchy Levels	Product	Not in scope, represented by data
	Field Device	Not in scope [41, 13], represented through edge AAS
	Control Device	ECS runtime [41] with deployed services, in particular Resources and Monitoring Layer with contributions from upper layers
	Station	ECS runtime [41], possibly with access to more powerful resources or UI capabilities for executing or controlling deployed services. Includes Resources and Monitoring Layer with contributions from upper layers
	Work Centers	Reusable Intelligent Services Layer, in particular Data Integration component
	Enterprise	Application Layer
	Connected World	On top of Application Layer via Applications AAS, including connected IIP-Ecosphere platforms
Life Cycle Value Stream	Type	Component and AAS types prescribing structures
	Instance	Deployed component and AAS instances

In term of the Industrial Internet Reference Architecture [23], this document can further be understood as a continuation of the usage view(point) [41], the functional view [13] In terms of a platform architecture as well as its implementation.

3.1.2 Stream (Data) Processing

In an IIoT/Industry 4.0 setting, often the processing of data is viewed in terms of streams of data items (or tuples), e.g., produced in regular fashion by a machine, taken up by edge devices for pre-processing, protocol transformation or retro-fitting, handled further by other devices and (partially) stored in some data stores, e.g., time series data bases. In contrast to other forms of data processing, e.g., batch processing, data stream processing can fulfill (soft) realtime requirements, of course, depending on the (relative) speed of the individual data processors.




Figure 4: Viewing IIoT and Industry 4.0 as data streams.

Figure 4 illustrates the basic components of such a stream processing approach, considering “the machine” on the left side as constant (conceptually endless) data source. The data produced by the machine is taken up by a data transformer (e.g., preprocessing, anonymization), passed to a second transformer (e.g., artificial intelligence) and finally to a sink (e.g., data store, dashboard). From a different point of view, the data flows in forward manner from source to sink. The edges in such a graph indicate the data flow and the nodes the data processors. There could be more processors,

different kinds of processors or more complicated forward flows that we do not touch in this brief introduction. Please note, that there is no need for synchronous processing in the nodes, in particular in the transformers. With synchronous processing, we mean that a transformer operates like a function in mathematics, i.e., for an input tuple it produces in the same step an output tuple. In contrast, asynchronous means that the processor receives data item(s) and at some point in time later it may emit any number of tuples (including none at all).

In Figure 4, there are also two horizontal lines, indicating borders of physical devices, e.g., the first two streaming components could be running on an edge device, the second transformer on a further device, and the sink on a third device, e.g., a central server. The distribution of components is not fixed, e.g., depending on resource usage, the second transformer could also be executed on the first device or the first transformer on the second device.

Several approaches to stream processing rely on untyped data, i.e., the transformer implementation decides based on the available data fields, what to process. Such an approach can easily fail at runtime, when processing nodes are combined that cannot work together, with negative outcomes ranging from loss of data to runtime errors or exceptions. In contrast, we rely on typed data flows, i.e., for each forwarding edge the type of data item(s) is known during design and built into the system. As the design of data processors and data flows will be captured in the configuration model, checking for type and streaming compliance before realizing or instantiating the system becomes possible. In Figure 4, the forward flow indicates three data types, t_1 , t_2 and t_3 . Please note that depending on the requirements and the design of the data processing, the types may be the same or they could differ, e.g., indicating that a processor adds or removes data fields.

While in many applications, a forward flow is sufficient, in particular in IIoT/Industry 4.0 settings it could be desirable, that an upstream processor shall send back data to a downstream processor, e.g., a decision node after one or multiple artificial intelligence nodes shall inform the machine at the data source that some processing parameters must be changed. Akin to the forward flow, we allow for backward flows. It is just a matter of modelling convenience that we define the forward flow in terms of nodes and connecting edges, while we consider the backward flow as typed notification data channel (t_4) of a sender and potentially multiple receivers.

3.1.3 Asset Administration Shells

The IIP-Ecosphere platform aims at complying with, integrating of and extending existing standards and technologies in I4.0 (R7, R14). This applies to protocols, formats but also model standards such as the Asset Administration Shells (AAS). For short and without aiming for a complete description, an AAS is an information model, which aims at modelling a physical or virtual Asset in terms of an asset specification as well as nested, detailing sub-models. Sub-models may consist of typed properties, operations and heterogeneous collections of sub-model elements. AAS and sub-models can be classified as static (all information is determined when creating the AAS), dynamic (some information may change at runtime) or active (callable operations are provided). Similarly, properties and operations can be static or dynamic, whereby in the dynamic case both element types can be linked to an implementation, e.g., provided by a remote implementation server, and thus change value (access) or implementation over time. In particular, AAS for different assets of different vendors can be provided, linked and integrated, e.g., to link the AAS of a device utilized by the IIP-Ecosphere platform into the platform AAS in order to have, e.g., the digital nameplate (for industrial equipment [43, 2]) or the documentation of the device at hands. Moreover, composite AAS can be created, representing, e.g., a complex machine consisting of AAS of the utilized components.

According to the requirements (R7), the IIP-Ecosphere platform shall describe all (distributable) components, interfaces, functions and deployment targets in terms of AAS. Thus, each of the components of the platform that forms an individual asset (of a certain vendor) shall receive an own

AAS (as indicated in Figure 1). Moreover, the platform itself shall provide an own AAS and each of the discussed layers shall provide one or more sub-models to link the layers against each other (whereby the sub-models may and shall link to the vendor AAS of the individual assets, e.g., edge devices). As far as feasible, the IIP-Ecosphere platform will utilize existing approaches and standards to define the AAS, but also define own ones where needed, e.g., to characterize the capabilities of deployment targets such as edge, server or cloud devices [41]. We will detail the platform AAS and its structure in Section 5.

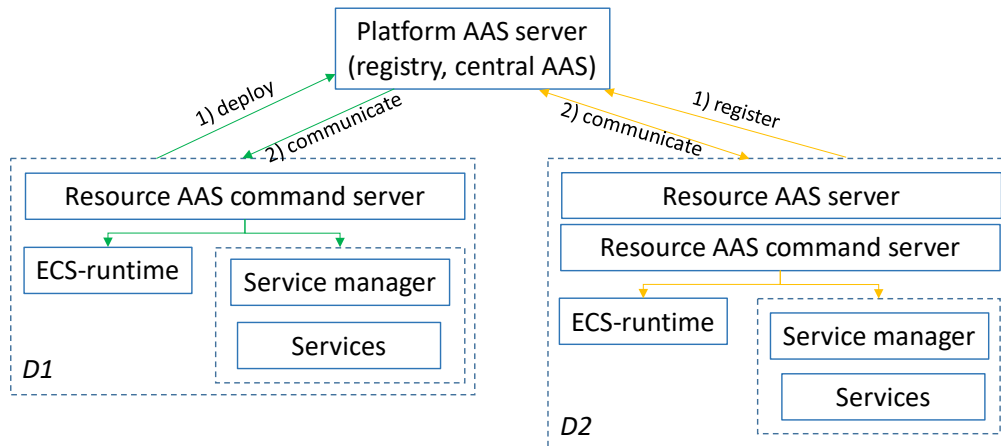


Figure 5: AAS deployment options (D1 remote deployment, D2 local deployment)

As typically several distributed compute resources are involved in a platform installation and each compute resource shall be described with an own AAS (model, sub-model or as part of joint model/sub-model), it is helpful to introduce now two basic AAS and component deployment patterns. Figure 5 illustrates the central IT side (the “Platform AAS server”) and two distributed resources *D1* and *D2*, e.g., edge devices. An AAS can be served locally and only be registered in a central registry or it can be deployed remotely to a central server. Serving an AAS locally requires a related web server process (“Resource AAS server” in *D2*), i.e., a further process to be executed on a resource. Deploying an AAS centrally avoids such local server processes, but may lead to increased communication with the central server and, in the case of dynamic or active AAS that allow for dynamic properties and operation calls, also to redirections of requests via the central server to the resource. To handle requests of dynamic or active AAS, the resource must run a (further) server instance, the “Resource AAS command server”. A similar server process must exist on the central IT side of the Platform AAS server to offer dynamic properties or operations. In the resource case, this “Resource AAS command server” may forward operations to further processes, or, if the processes are already known when the resource AAS is constructed, also specific server processes, e.g., for the service manager running in an own container, can be linked to the AAS and directly contacted to serve AAS requests.

3.1.4 Component Interaction Overview

In the previous sections, we introduced the layers and the high-level components of the platform as well as the basic concepts of AAS. In this section, we provide a brief overview on the component interactions for a basic walk-through of platform operations. The individual sections on the components in Sections 3.5-3.14 will provide more detail on the interactions. In addition, Section 3.15 will address the cross-cutting topic of testing support for services and applications.

The aim of this walk-through is to bring up the ECS runtime, the service manager (in terms of a container), some services, to let the services run and to stop all parts in reverse order. Services are described in terms of a service mesh forming individual applications (we will detail how to define such a mech in Section 6). The required high-level interactions are illustrated in the sequence diagram in Figure 6. We will go through them now from top/left to bottom/right.

1. At the beginning, the platform AAS-Server is running. An ECS runtime is started for a certain resource, e.g., an edge device. The ECS runtime instance then deploys its own sub-model characterizing the device with container operations and a collection of available containers (initially empty) into the platform AAS (remote deployment). A scheduled background process of the ECS runtime is started to inform the platform AAS about the actual resource state (resource monitoring, not shown in Figure 6). Depending on the device, the ECS runtime may provide information about an existing device AAS or create a device AAS on its own (one particular point of openness as the device vendor may or may not provide an AAS). This information is linked from the platform AAS.
2. Via the user interface (UI), the user requests a list of available resources. The UI reads out the AAS submodel for resources including the ECS runtime instance started in step 1 and prints out device information including the actual resource usage. In a similar manner, further information can be obtained, e.g., the available services, the defined applications, the packaged service artifacts or the available containers.

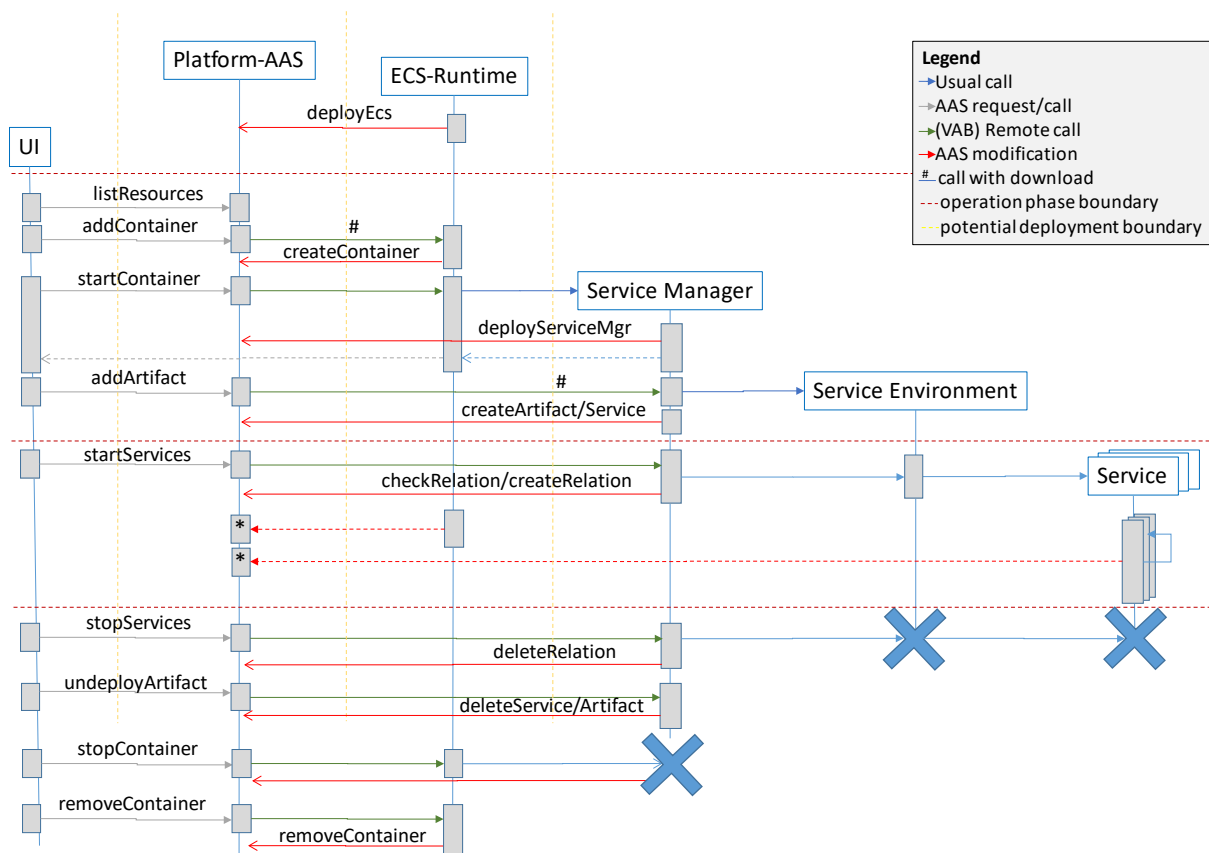


Figure 6: High-level component interaction for basic platform interactions.

3. The UI requests adding a container via the ECS operations known to the platform AAS, leading to a remote method call to the ECS runtime (AAS implementation server). For this walk-through, we assume that the container contains the service manager and provides the technical dependencies for services to be executed on the respective device. Starting a container may lead to a download of the container from a central platform server (indicated by #) or from a file system of the device. Information about the container instance is made available to the platform AAS by creating a structure in the containers submodel of the platform AAS.
4. The user requests starting the container added in step 3, i.e., the UI calls the respective platform AAS operation, leading to a remote call to the ECS runtime (AAS implementation server), respective operations in the container management implementation, e.g., Docker,

and, ultimately, when the container is running, to an automated start of the service manager. In turn, the service manager deploys information about itself, e.g., service operations, into the platform AAS, more precisely into the device entry created by the ECS runtime so that services on the underlying device can be managed.

5. So far, no service is known. The user requests to add a service via an operation of the platform AAS, leading to a remote method invocation to the Service Manager (AAS implementation server). In turn, as for the container, the service manager may download an implementation artifact containing the service execution environment and the individual services for the actual device. The service manager adds entries for the artifact and all contained services to the respective sub-models of the platform AAS.
6. The user requests the start of all services for the device addressed in the steps above. The service manager starts the service environment and creates the service instances in the sequence of dependencies, i.e., starting with the service having no data dependencies or for which all prerequisite services are already running. During this step, several free network ports may be acquired for internal communication, relations to a global or a local protocol server/broker may be established and individual operating system processes for the services may be started. Further, connector settings may be adjusted to device-provided services as stated in the device AAS, e.g., specific IP address, actual port or even the startup-time selection of the right protocol implementation depending on the device-specified protocol version, e.g., MQTT v3 vs. MQTT v5. These detailed technical procedures are not shown in Figure 6. During service startup, the Service Manager checks the service relations in the platform AAS (services sub-model) for service availability and, as soon as the service is up, creates a relation entry linking two subsequent services in the service mesh of an IIoT application running on the platform.
7. The services are running now, receiving data via the machine/platform connectors, executing functionality specified for the actual application, e.g., AI-based inference. During the execution, background processes collect data for the device and the individual services and inform the platform AAS about changing runtime states, e.g., resource consumption. Here we also indicate in Figure 6 the resource monitoring of the resource mentioned in step 1.
8. The user requests to stop the running services via a respective operation of the platform AAS, which causes a remote method invocation to the Service Manager. In turn, the Service Manager removes the service relations in the platform AAS and stops the service environment and the services.
9. The user decides that the artifact will not be used any longer, i.e., a platform AAS operation is called and causes a remote method call to the Service Manager, which removes service and artifact entries from the platform AAS.
10. As also the service management container shall not be used anymore, a command from the UI to the respective AAS operation leads to a remote method call to the ECS runtime, which commands a stop of the container through the underlying container implementation.
11. Ultimately, the container shall also be removed from the management realm of the device, leading to a further remote method call to the ECS runtime, performing a removal of the container information from the platform AAS.

The horizontal dashed, red lines in Figure 6 indicate phases of the operations, i.e., startup (step 1), preparation of containers and services (steps 2-5), service operation (steps 6-7), shutdown (steps 8-11). The vertical yellow dashed lines indicate a potential distribution to different logical or physical devices. Extreme cases are that all components run on the same device, e.g., for testing, or that UI, platform AAS, ECS runtime and service manager/services are installed/started in separate devices.

It is important to emphasize that the “user” in this walk-through may be a human, an automated process of the UI or the platform itself acting on behalf of the user. An automated process could be a deployment script, which lists the assignment of containers and services to resources so that the UI can execute the desired deployment in a single step (broken down into the actions discussed above). Also, as intended for the Deployment component, the platform may calculate such a distribution of services and containers and enact the deployment automatically or autonomously.

Moreover, it is important to mention that not all services of an application must be started at once by a single command. If services shall be distributed to different devices, individual start commands are required to enact this form of distribution. However, for a human, defining such a distributed deployment stepwise is cumbersome. For this purpose, the platform offers mechanisms to script such deployments or – as envisioned for future versions of the platform – to calculate and enact them automatically.

3.1.5 Virtual Character of the Platform

As stated in Section 1, the IIP-Ecosphere platform shall be designed as a **virtual platform** (R3), i.e., a platform that offers services on top of existing already installed platform functionality. The idea is that the Connectors component in the Transport and Connection Layer map relevant underlying platform information and functionality into the IIP-Ecosphere platform. Where feasible, this mapping shall happen in the form of AAS as it allows for an overarching information model, but also further approaches like OPC UA or MQTT may be used. We see here three alternatives, focusing on AAS as the default approach, potentially using a transport protocol like MQTT for high-speed data connections:

1. An underlying platform provides its own AAS and manages the access to selected functionality and data. Theoretically, this AAS could be mapped side-by-side into the AAS of the IIP-Ecosphere platform. Then, layers such as deployment device management, or monitoring could directly utilize the information. Therefore, a standardized AAS structure for manufacturing platforms would be desirable, but such a standard currently does not exist.
2. The AAS connector of the IIP-Ecosphere platform can map the AAS of the underlying platform into the format of IIP-Ecosphere. Of course, this adds additional overhead and in some cases a mapping may not be possible at all.
3. One of the other IIP-Ecosphere connectors provides a protocol that allows mapping the underlying platform and its operations into the IIP-Ecosphere AAS format. This approach may require manual programming, while the second approach might be realized easier through mapping and code generation.

Besides having the AAS of an underlying platform, relevant components of the IIP-Ecosphere platform, in particular the resource management and monitoring component are required to operate with multiple AAS instances (for now based on the IIP-Ecosphere AAS structure).

3.2 Development Streams

Realizing the IIP-Ecosphere platform in one big shot is not realistic. As already indicated in the previous sections, we rely on incremental and agile development, so that delayed concepts, designs and results can be integrated flexibly, e.g., after initial experimentation with the available platform components. For the increments, we identified three development streams as indicated in Figure 7.

In the first stream, we aim at the basic functions, i.e., support for creating AAS, data transport, data connectors, basic service interfaces and management as well as the environment for running services on edge, cloud or server installations (ECS runtime). In agile spirit, these realizations must be functional and tested but not complete, e.g., it is more important to start/stop dependent services/containers rather than to perform a runtime migration of services or containers.

In the second stream, advanced functions are added and functionality missing from the first development stream may be realized. At latest, missing functionality will be integrated with the improved and advanced functions in the third development stream.

We do not indicate a detailed time schedule for the streams or functionalities here. The first development stream was completed in Spring 2021, the first release of the second stream was made available in Summer 2021 (along with this version of the handbook) while the second development stream was mostly completed until end of 2022. Each stream is manifested by at least one release of the platform. This version of the platform handbook summarizes the state for the second release in the second development stream. Although this version adds several new features (dynamic service ensembles, data processing library, container mocking, service test generation, use of semantic IDs in AAS, distributed testing and evaluation environment), it mainly targets the revision and stabilization of functions (asynchronous streams, monitoring, web-based UI) as identified during the development of the Hannover Messe 2022 demonstrator [11]. Many features for the second development stream (initial advanced functions) are completed. For security and data protection, a platform service for anonymization is integrated, while data lakes/stores are still in conception. Similarly, the data integration and the AI toolkit partially do exist (and their conceptual foundations are still in progress) but are currently not (completely) integrated. Adaptive services, e.g., runtime parameterization or switching between compatible service alternatives was already introduced with version 0.3.0 of the platform.

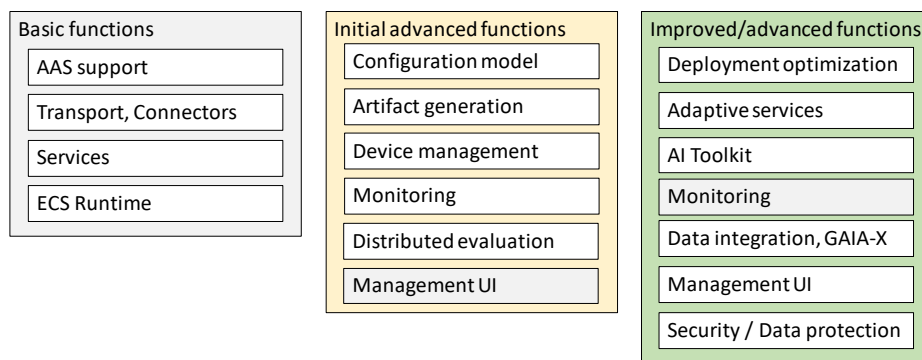


Figure 7: Development streams for the IIP-Ecosphere platform

3.3 Overall Requirements

In general, all platform layers and components discussed below must take the following general requirements from [13] into account:

Table 2: General platform requirements in [13]

Requirement	Summary
R1	Vendor and technology neutral platform
R2	Use of standards
R3	Design as a virtual platform
R4	Design based on components and services
R5	Use of Open Source, with respect to the licensing rules of IIP-Ecosphere
R6	Open for optional/commercial components
R7	Use of AAS for interfaces
R8	Use of systematic variant management techniques
R9	Means for availability
R10	Soft realtime processing (<100 ms) for production-critical functions
R11	Documentation (also in terms of this handbook)
R12	Documentation of processing steps (of applications, supporting data privacy)

As already indicated in Table 2, [13] also specifies quality requirements such as R10. Besides security and data protection requirements, there are also data frequency and volume requirements that are not so obvious, in particular as they are assigned to specific topics/components of the architecture in [13]. To provide an overview, we discuss them here on a global level for the entire platform.

In Table 3, we summarize the cross-cutting quality requirements, i.e., in particular those that may require specific considerations regarding time-critical functionality such as the (stream) processing or data transport. Although the IIP-Ecosphere platform aims at the deployment of components to edge devices, both, the services as well as the platform operations there belong to the IT realm so that OT requirements such as R35 or the OT sensor sampling frequency mentioned in R28 do not directly apply. However, a machine pulse of 8 ms (R28) as well as an hourly throughput of 7 GByte as well as an expected size of data items with 50 values (R19a) are highly relevant for judging the performance of the IIP-Ecosphere platform. As also mentioned in [13], not all data volume and frequency requirements were indicated while collecting the requirements from the partners, i.e., the platform shall aim for even higher speed (such as a 50 ms cycle time) or a throughput of 600 GByte per day.

It is also important to recall from [13], that the IIP-Ecosphere platform is primarily responsible for its mechanisms and included services, i.e., providers for services to be packaged with the platform will have to obey the quality requirements in [13] and in particular Table 5. Further, as also discussed in [13], the platform is not responsible for the quality of external services, e.g., application-specific or user-specific services (while measures may apply to report or terminate services that potentially taint given runtime requirements).

Table 3: Overview of (global) quality requirements on data frequency and volume

Requirement	Summary
R10	Soft realtime, response time < 100 ms for production critical functionality
R19a	Sample data set of 50 values of different types all 20-30 s
R19e	Output data shall be provided all 5 s
R21	Low impact on data throughput
R22	Overall platform throughput of 500 GByte per year
R28	OT sensor sampling frequency 0.2 ms, machine pulse 8 ms, step pulse 5 s, process pulse 25 s (mentioned in the explanation of the cloud requirement R28)
R35	OT sampling frequency of 2 ms
R91	7 GByte per hour as input for data integration, which may be aggregated to 2 Gbyte per hour.

As an illustration, we discuss the quality requirements now in terms of hypothetical numbers. From the data transport perspective, the requirements command that each machine can ingest a data item with around 50 values each 8 ms, i.e., 125 messages per second. This leads to at least 450.000 messages per hour (per machine/edge device). If we assume a size of 654 Byte payload (actual size of a simple JSON serialization of such as message), a source produces around 280 Mbyte per hour (just focusing on the raw data payload, i.e., not on additional information, e.g., for routing or meta-information as stated in R79). On a platform-level (R91, R22), aggregating components of the IIP-Ecosphere platform will have to cope with multiple parallel streams of this kind, which requires 26 such streams to reach the requested 7 Gbyte (in a real setting with payload and overhead). Of course, the distribution may be different, i.e., more streams at lower ingestion frequency or less streams at maximum frequency, potentially with image payloads, to reach several hundreds of GBytes per hour.

In the discussion of the individual layers/components, we will refer to these general requirements and re-iterate the argumentation only for affected layers or layers that already have been (initially) evaluated.

3.4 UML Profiles

The IIP-Ecosphere architecture model is based on three UML profiles, the IIP-Ecosphere profile introduced in Section 3.4.1, the UMLsec [24] profile for security modeling in Section 3.4.2 as well as the security and privacy profile in Section 3.4.3. All three profiles aim at classifying and defining orthogonal information to be attached to individual modeling elements. While the IIP-Ecosphere profile as well as the security and privacy profile are mostly of descriptive nature, i.e., indicate additional information such as open source licenses and component versions, the UMLsec profile is the basis for automated security analyses of UML models using the CHARISMA tool.

3.4.1 IIP-Ecosphere Profile

The aim of the IIP-Ecosphere profile is to classify and categorize modeling elements in the IIP-Ecosphere architecture, i.e., to express orthogonal semantics in a uniform manner. We will now briefly discuss the individual concepts and parts of the profile.




Figure 8: AAS stereotypes in the IIP-Ecosphere profile (comments cropped).

One cornerstone of the IIP-Ecosphere platform is the exploration and use of **Asset Administration Shells** (AAS, R7 in [13]). The partners decided to use AAS in particular to describe interfaces of the platform (internal, external) and communication with these interfaces in a standard-based uniform approach³⁶. Thus, from an architectural point of view, it is relevant to model (structural) AAS aspects. We use AAS in terms of classes, interfaces and operations tagged by the stereotypes depicted in Figure 8. A class/interface can be marked by the «AAS» stereotype to express that there shall be an AAS providing access to the contained data. An «AAS-DataConnector» is a communication endpoint³⁷, e.g., for soft-realtime (streaming) connections. Such endpoints that are currently not part of the AAS standard³⁸. An «AAS-Property» is a static or dynamic attribute of an AAS. UML properties may also indicate that a substructure (i.e., an «AAS-SubModel») shall be exhibited by an AAS. Moreover, AAS may describe functional interfaces using the «AAS-Operation» stereotype.

Moreover, an «AAS-Client» is per se not an AAS element. In the IIP-Ecosphere platform, an AAS-Client is a supporting class implementing how to access properties or how to execute operations. These classes shall be defined along with the respective AAS and can be tested directly against the AAS.

As our approach to modeling is pragmatic and agile, we do not aim at covering all possible aspects of AAS. Please note that the stereotypes just indicate that the respective information shall be represented in a realizing AAS. We do neither model the concrete names used in a realizing AAS, the completeness of models or sub-models nor any sequence of contained AAS elements. Besides properties that can

³⁶ Design guidelines for AAS must still be agreed upon by the partners or discussed with other projects. Although this affects the implementation, the actual AAS design guidelines are outside the scope of the architecture, i.e., we focus here just the relevant aspects such as properties, operations and links/dependencies.

³⁷ Following [41], we do not use the term “endpoint” in this document rather than “data connector”. For links among data connectors and endpoints, we use the terms “relation” [41], “data flow” or “data path” [13].

³⁸ There is ongoing work on standardizing communication endpoints. If possible, we will adopt this upcoming standard in a later release.

change their value at runtime or sub-models that occur on demand, dynamic relations among AAS elements can be modeled by dependencies marked with the AAS-dynamic stereotype.

A second cornerstone are **services**, in particular to encapsulate platform functionality or (re-usable) AI methods. According to the profile (Figure 9), a service can be modeled as an interface (with implementation aspects hidden) or as a class (i.e., a namespace with properties and operations). Related to services are (platform) **connectors** that ingest data into the platform or are involved in offloading data/processing to other platforms or a cloud. Typically, for one connector type multiple alternatives are offered and also additional connectors can be added (openness, e.g., R14, R16). To indicate these elements, the profile contains a generic «Connector» stereotype that can also be used to indicate Cloud connectors and (for security) optional IDS connectors.




Figure 9: Service and connector stereotypes in the IIP-Ecosphere profile (comments cropped).

In IIP-Ecosphere, services shall be deployed in terms of virtualized **containers**. Therefore, implementing elements can be marked as «Container» (Figure 10). Further, besides services, individual platform components can be marked as «Distributable» while parts not marked as «Distributable» shall remain part of a central platform installation.




Figure 10: Container and distribution stereotypes in the IIP-Ecosphere profile (comments cropped).

To simplify the models, i.e., to avoid repetitively modeling of typical mechanisms or collaborations, the profile allows indicating **architecture, design or implementation patterns**³⁹, to explain/detail a model element in an uniform manner, but also to guideline the implementation. Figure 11 depicts simple patterns like architectural layers⁴⁰, delegation of control to another element via an association, read-only attributes (without corresponding setter)⁴¹, builder pattern⁴² (or classes that shall use this pattern to realize read-only attributes) or visitor pattern⁴³. Figure 12 illustrates stereotypes for marking an object factory⁴⁴ (an exchangeable mechanism that creates instances) or plugins as one form to extend platform functionality at defined points. A «Plugin» is detailed by a «PluginType» providing more

³⁹ An important reference here is the GoF book [17], but for simplifying the understanding, we just provide some Web references.

⁴⁰ https://en.wikipedia.org/wiki/Multitier_architecture

⁴¹ UML and Papyrus offer a read-only meta-property of the meta-class Property. However, displaying this information in the diagram is tedious, so we just define the corresponding stereotype read-only.

⁴² https://en.wikipedia.org/wiki/Builder_pattern

⁴³ https://en.wikipedia.org/wiki/Visitor_pattern

⁴⁴ https://en.wikipedia.org/wiki/Factory_method_pattern

information on how to implement/register the plugin. The default type is JSL, the Java Service Loader⁴⁵, a simple mechanism on an implementation to its (descriptor) interface without direct dependencies in code. These patterns support the openness of the platform, e.g. extensibility for optional components in R6 [13].




Figure 11: Basic architecture/implementation patterns in the IIP-Ecosphere profile (comments cropped).

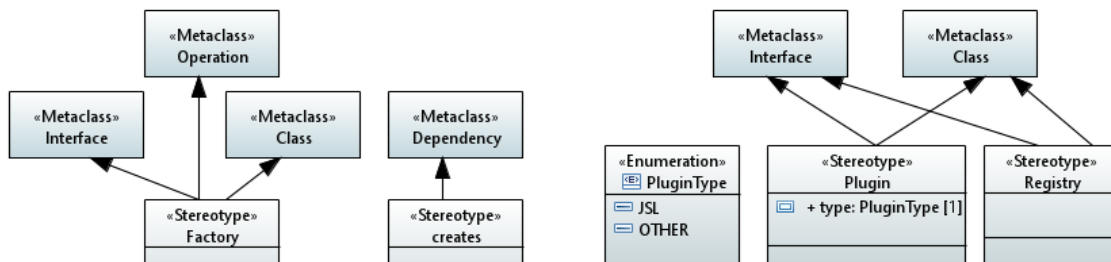


Figure 12: Factory and plugin/registration patterns in the IIP-Ecosphere profile (comments cropped).

Due to the AI nature of IIP-Ecosphere it is important to recognize that code written in various **programming languages and under licenses** must be integrated (R5, R6, R113 in [13]). The stereotypes in Figure 13 allow indicating these two dimensions also to locate potential pitfalls. External components not marked with the OpenSource stereotype are meant to be commercial/proprietary and shall be only used as alternatives rather than mandatory or default components, i.e., their installation shall be left to the user's choice. «OpenSource» components are characterized by their license(s) and their version. If no version is indicated, no decision was made so far, i.e., the component was not integrated so far.

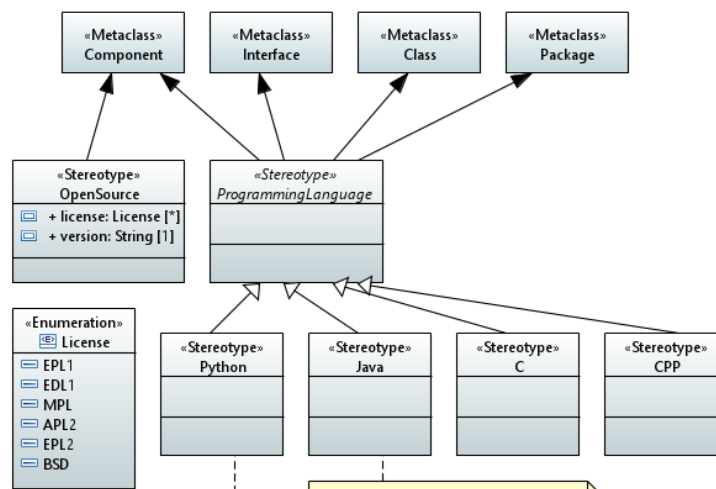


Figure 13: Licenses and programming languages in the IIP-Ecosphere profile (comments cropped).

Within the architecture model, it is also relevant to mark the **maturity status** of individual parts, e.g., to distinguish initial models from detailed models that are actually implemented (Figure 14). Some parts (within models) may not have been realized so far and can be marked with «Omitted». The

⁴⁵ <https://docs.oracle.com/javase/9/docs/api/java/util/ServiceLoader.html>

maturity status can be attached to a model or individual modeling elements if applicable, but in particular also to comments. We use comments to document the aim/contents of each UML diagram (often cropped in this document) and may then attach the maturity stereotype to that comment.

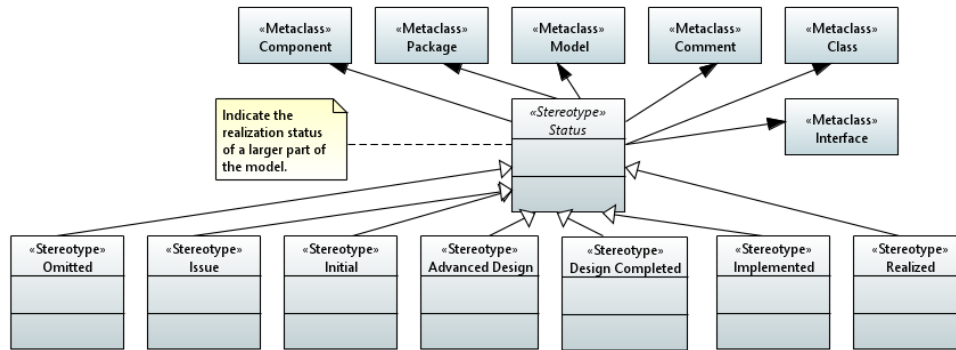


Figure 14: Maturity status for comments, packages or models.

Another special contribution of the IIP-Ecosphere platform is to aim for an encompassing and consistent **configuration model** that ranges from devices over services to service orchestration and covers static pre-runtime and runtime aspects, e.g., for deployment optimization or self-adaptation (R120-R126, R107 in [13] as listed in Section 3.1). We use the Integrated Variability Modeling Language (IVML) [10] to describe the configuration model and we indicate aspects of the variability modeling in the architecture models. IVML is realized in terms of the EASy-Producer [36] toolset, an external open source component that we integrate into the IIP-Ecosphere Platform. For short, the IVML configuration meta-model of IIP-Ecosphere (represented as information items marked with the IVML stereotype shown in Figure 15) defines the structure, configuration options and validity criteria for all potential platform instances. The configuration (also an IVML model) instantiates the meta-model and details the configuration decisions for a specific platform instance, e.g., on which server the platform AAS will be located, which concrete services are available etc. One particular architectural aspect is the structure of the IVML (meta-)model and its relation to the layers of the platform. The (meta- and configuration) model consists of individual modules (called projects). We specify this decomposition of the configuration model into modules (represented as information items tagged with «IVML») in terms of dependencies decorated with «IVML-Import». Ultimately, mechanisms of EASy-Producer will validate the configuration and automatically modify, include, exclude, generate or package artifacts that finally make up the configured platform instance.

A further architecturally important aspect are the components and classes realizing the variations defined in the configuration model (i.e., the implementation parts to be included, excluded, modified, linked with glue code etc.). We use the Software Product Line [42] notion of «Alternative» or «Optional» artifacts and mark the respective components using the stereotypes shown in Figure 15. Examples are alternative transport connectors (one must be selected) or optional components (that can be part of the platform instance or not) such as IDS or cloud connectors. As these stereotypes are intended to be illustrative and explanatory rather than for defining a configuration or an artifact model, we do not include further variability details as done in typical variability profiles, e.g. in [44, 18]. In particular, components marked with «Alternative» or Optional shall be designed and implemented carefully with respect to their dependencies, i.e., leaving out an «Optional» component or replacing an «Alternative» component must not render a platform instantiation invalid unless the governing configuration is invalid.

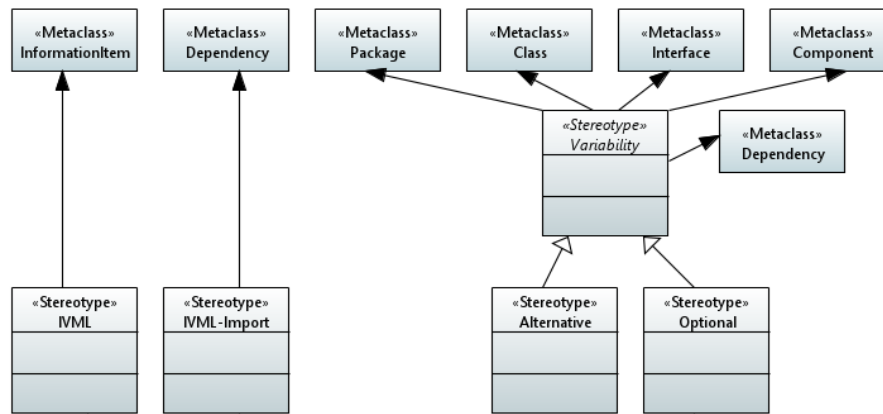


Figure 15: Configuration modeling and variability management stereotypes (comments cropped).

One form of instantiating a configuration model is to **generate code**, e.g., data transport code or glue/binding code between a hand-crafted implementation and platform interfaces to ease the development of consistent applications. To indicate that parts of the architecture are intentionally left open as they will be filled through generation from the platform configuration model during platform instantiation, we mark these parts by the «Generated» stereotype shown in Figure 16.

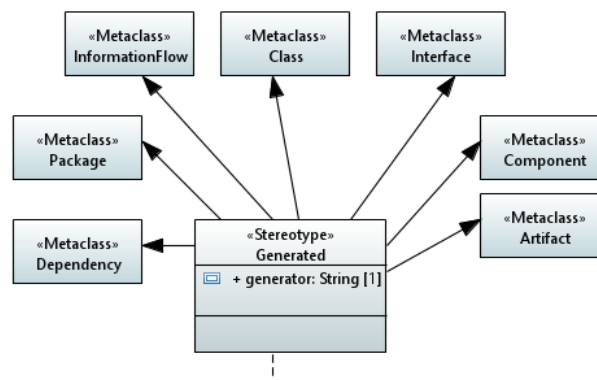


Figure 16: Stereotype for generated code (comments cropped).

The IIP-Ecosphere platform shall provide **self-adaptive capabilities at runtime** based on the (runtime part of the) configuration model (see R102-R109 and R120-R126 in [13]). Examples of platform elements that could be adapted are container deployment locations or actual AI services used in the same deployment location. To indicate model elements that are related to adaptation, the IIP-Ecosphere profile defines the stereotype «Adaptation» (Figure 17). One particular example for the application of this stereotype is to mark states in a state diagram that would not be needed without self-adaptation functionality, but which are required for self-adaptation.

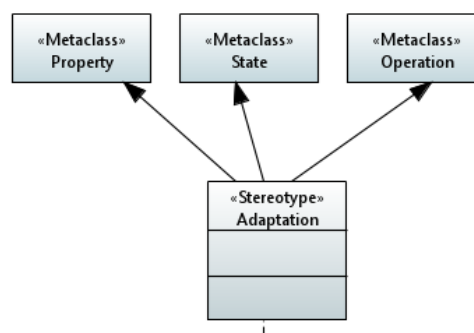


Figure 17: Marking model elements as support for self-adaptation.

In some situations, we also include **experimental** components in the architecture, in particular to introduce a certain concept that is used later on in the discussion, e.g., in a validation. To indicate such experimental components, we utilize the stereotype shown in Figure 18. For various reasons, experimental components may not be found in the IIP-Ecosphere source code repository and they may also not be subject to the continuous integration.




Figure 18: Marking experimental components

In Papyrus, it is possible to define a stylesheet to adapt the formatting of modeling elements based on the applied stereotypes. We will use this mechanism to mark important stereotypes, e.g., issue comments or omitted elements, in a uniform manner such as uniform fill or text colors.

3.4.2 UMLSec Profile

UMLsec⁴⁶ provides a model-based approach to develop and analyze security critical-software, in which security requirements such as confidentiality, integrity, and availability are expressed within UML diagrams. The UMLsec language is provided as a UML profile and can be imported into existing UML tools. In UMLsec, different stereotypes and tags are used to annotate UML diagrams with security properties. UMLsec provides various security checks to ensure the annotated properties. The CARISMA tool performs the corresponding security checks. The idea of UMLsec is to provide maximal analysis power while allowing to use everyday development tools for the development process.

While the UMLsec profile is defined as a light-weight UML extension, it is also possible to define it using heavyweight extensions to specify the change of semantics. One can make use of an extended meta-model (analysis model). This analysis model provides the possibility of more complex analysis by extending the basic UML meta-model.

As mentioned above, UMLsec provides different security checks to verify whether a security property in a system is violated, and a security mechanism is needed to restore it. In this document, we explain two security checks, namely *secure links* and *secure dependency*. *Secure links* is used for the description and the analysis of secure data flows over connections between the artifacts in a UML deployment diagram, which describes the physical layer of a system. *Secure dependency* ensures that various dependencies between interfaces in a structure of a system model respect the security requirements of the data communicated across them.

3.4.2.1 Secure Links Check

The physical layer of a system is modeled by a deployment diagram, including physical nodes, the communications between them (modeled by links), the (software) artifacts and the dependencies between the artifacts. The *secure links* annotation enables one to ensure the security of communications in a physical layer.

In UMLsec, to perform a security check, *adversary patterns* are required. Such patterns specify the potential access paths threatened by a certain attacker. Table 4.1, represents the *default* adversary, as

⁴⁶ <https://rgse.uni-koblenz.de/jj/umlsec/>

an example of an adversary pattern. For a given adversary of type A , the set $\text{Threat}_A(s)$ specifies which kinds of actions the adversary can apply to a node or a link marked with the stereotypes. For example, considering an unencrypted *internet* communication link, the default attacker ($\text{Threat}_{\text{default}}(\text{internet})$) can *delete*, *read* and *insert* messages transmitted over this link.

Stereotype s	$\text{Threat}_{\text{default}}(s)$
« <i>internet</i> »	{delete, read, insert}
« <i>encrypted</i> »	{delete}
« <i>LAN</i> »	\emptyset

Figure 19: The UMLsec default adversary pattern

The stereotype «*secure links*» implies the following conditions: for each dependency annotated with stereotype $s \in \{\text{«secrecy»}, \text{«integrity»}, \text{«high»}\}$ between two artifacts deployed on two nodes n, m , we have a communication link l between n and m with stereotype t such that:

- $s = \text{«high»}$, implies that $\text{threat}_A(t) = \emptyset$,
- $s = \text{«secrecy»}$, implies that $\text{read} \in \text{threat}_A(t)$, and
- $s = \text{«integrity»}$, implies that $\text{insert} \in \text{threat}_A(t)$.

For instance, if a communication link between two nodes n, m are annotated with «*internet*», and the dependency between two artifacts a_1 (deployed on node n) and a_2 (deployed on node m) are annotated with «*high*», then the security constraint associated with the stereotype «*secure links*» is violated: the dependency annotated with «*high*» demands that the set of threats of an adversary is empty, however, the communication link is annotated with «*internet*», meaning that the adversary is capable of reading, deleting, or inserting messages over the link between n and m . Consequently, the security requirement of the communications is not supported.

3.4.2.2 Secure Dependency Check

In UML, a dependency between two model elements is a relationship that denotes a model element requires other model elements for its specification or implementation. In other words, the complete semantics of the client element is either semantically or structurally dependent on the definition of the supplier element. The stereotype «*secure dependency*» implies that the security requirements have to be supported by both sides of the dependency (respective classifiers) and the dependency itself.

Later in this document (within the data security layer) we describe the UMLsec profile within the architecture model of the IIP-Ecosphere. Furthermore, we show how a CARISMA check can be performed on such models to verify the security level of the architecture models.

3.4.3 Security and Privacy Profile

To enhance the security and privacy of the platform, we create a dedicated Security and Privacy UML profile. The purpose of this UML Profile is to provide a catalog of security and privacy mechanisms to annotate corresponding UML models and the architecture of the platform. With such an annotation process, we can express appropriate mechanisms on a high abstraction level. In this way we give an overview of a complete security and privacy framework to the developers and designers of the system. Furthermore, this dedicated profile enables privacy and security audits. Audits lead to increased software quality.

The catalog represented by the profile introduces a means to structure privacy and security mechanisms in multiple abstraction levels. Namely, design strategies, sub strategies, patterns, and privacy enhancing technologies (PET) that can be applied to design of the platform.

The feature model in Figure 20 shows the four abstraction levels of the privacy and security concept. A feature model describes a set of features and their relations, here privacy and security features and their relations. The main structure of the feature model is hierarchical. But there are relations that do not follow the hierarchical structure. Some features may require other features. For example, Authorization requires Authentication. Other features may exclude each other, for example, Anonymity Set and Notify.

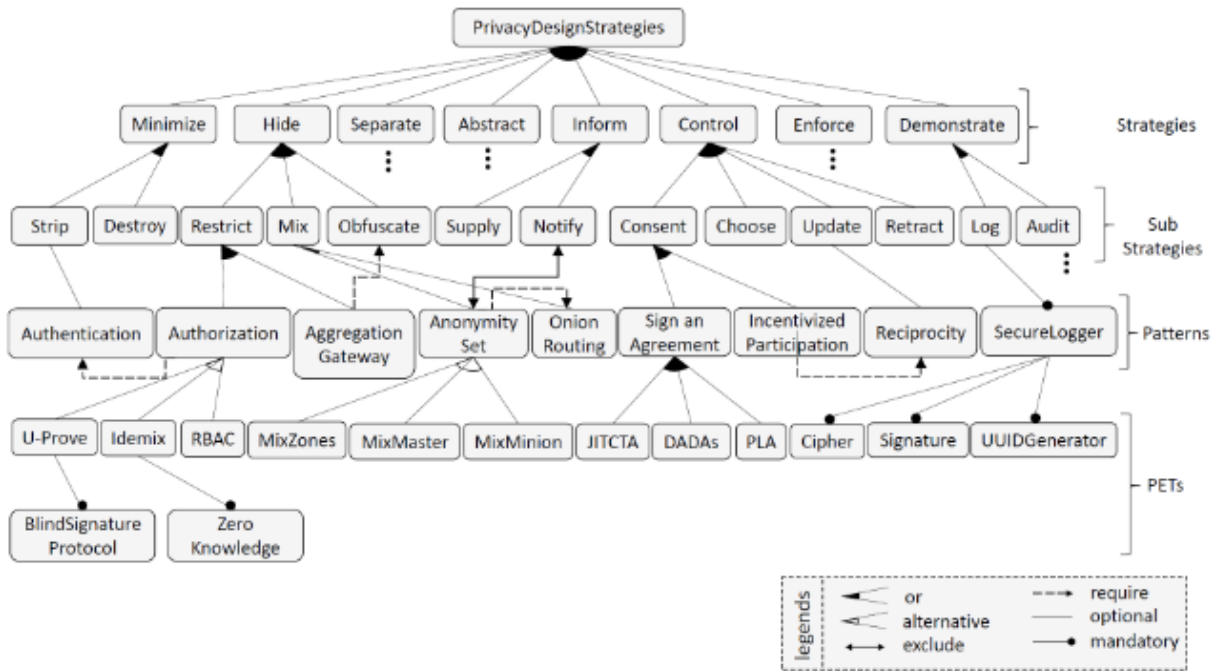


Figure 20: An excerpt of the feature model including privacy design strategies, sub-strategies, privacy patterns, and PETs (cf. [1] Figure 6.5).

Figure 21 shows the UML privacy and security profile we created based on the feature model shown in Figure 20. We adopted the hierarchical structure of the feature model and recreated it in terms of a UML profile. The profile has the same 5 levels as shown in the feature model, the root level PrivacySecurity, the Strategy level, Sub Strategy level, the Pattern level and the PET level.

For each level we defined to what elements in the model the stereotype can be annotated with that stereotype. For example, we can use the «Privacy» security stereotype to annotate components and packages. Stereotypes from the strategy level, like Hide, can additionally be used to annotate classes and interfaces.

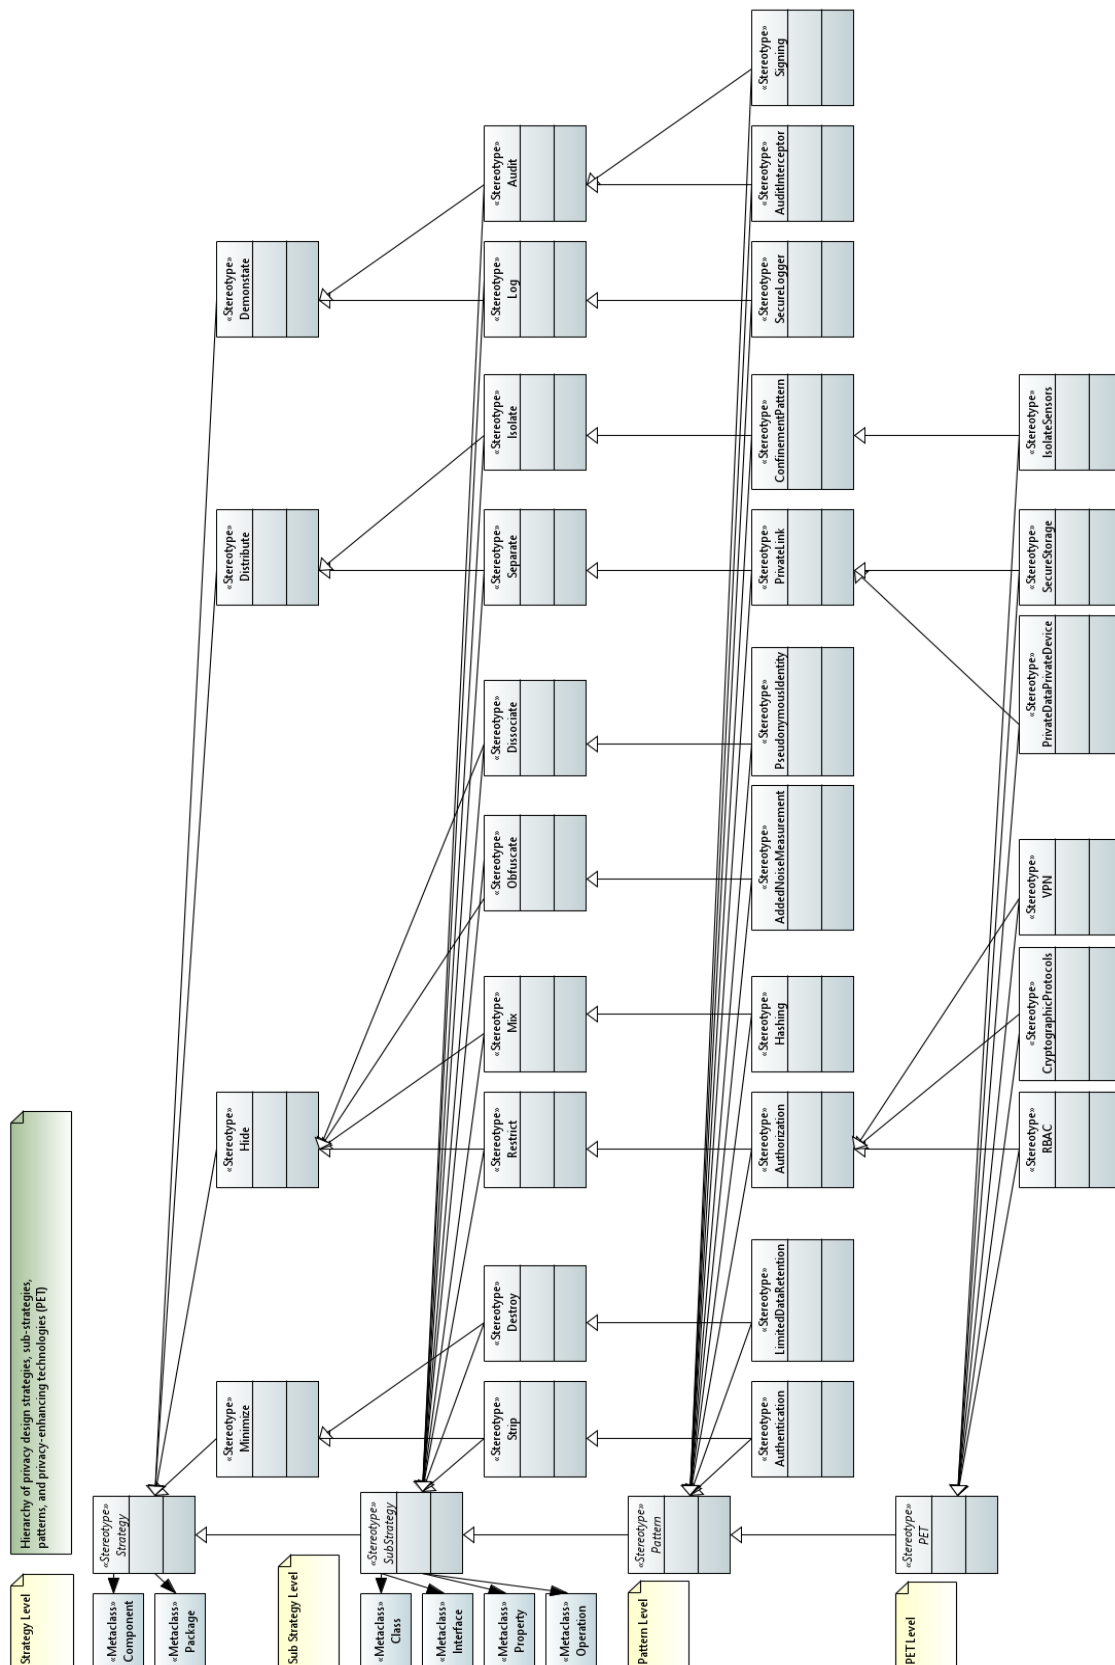


Figure 21: The Privacy and Security UML Profile (excerpt, cropped).

We discuss now specific examples on how the stereotypes from the UML privacy and security profile can be used. In IIP-Ecosphere, the profile can be used by business partners to communicate with each other about business secrets or to communicate with expensive production equipment. Unauthorized access to the system can cause severe damage to the companies using and trusting it.

We will now use the stereotypes to annotate our model with the role-based access control (RBAC) stereotype. In RBAC, the access rights are assigned to roles. Then individuals are assigned to the roles. This has multiple advantages over assigning roles directly to individuals. RBAC is a privacy enhancing technology. In our hierarchy, the RBAC PET is located in the Minimize Strategy, the Restrict Sub Strategy and the Authorization Pattern.

Figure 22 shows an interface (to be introduced in Section 3.6.3) annotated with the «Authorization» stereotype, and in that interface the write method is annotated with the RBAC and the Log stereotype.

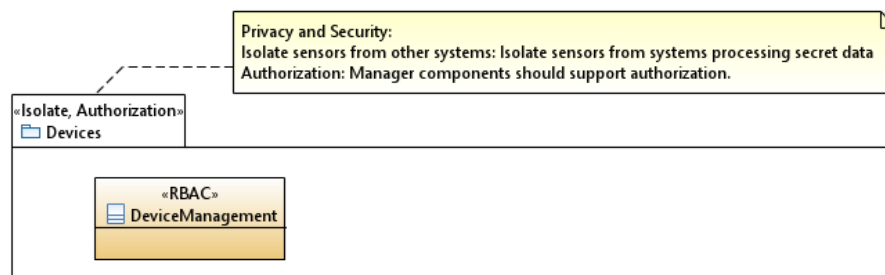


Figure 22: Interface annotated with Privacy and Security stereotypes

Figure 23 shows how the serialization package (to be introduced in Section 3.6.2) is annotated with the «Hashing» and «Signing» stereotypes. The contents of the package has been omitted in order to focus on the stereotype application. Serialization is one important part of storing, loading and transmitting data. With «Hashing» we can increase the integrity and with signing we can verify the origin of the data.




Figure 23: Package annotated with Privacy and Security stereotypes

Table 4 shows an excerpt of strategies, sub strategies, pattern and PETs that are suitable for the system. The design strategies, patterns, and privacy enhancing technologies are based on the work of Ahmadian [1]. The strategies are adapted from Hoepman [21].

Table 4: Design strategies, patterns, and privacy enhancing technologies for the IIP-Ecosphere architecture.
(cf. [1] Appendix F).

Strategy	Sub Strategy	Pattern	PET
Minimize	Strip	Authentication	
	Destroy	Limited Data Retention	
Hide	Restrict	Authorization	RBAC, Cryptographic Protocols, VPN
	Mix	Hashing	
	Obfuscate	Added Noise Measurement	
	Dissociate	Pseudonymous Identity	
Separate	Distribute	Private link	Private Data private Device, Secure Storage
	Isolate	Confinement Pattern	Isolate Sensors from System
Demonstrate	Audit	Audit interceptor	
		Signing	
	Log	Secure logger	

3.4.4 IoT Component Security and Privacy Profile

In parallel to the concepts provided in the previous section, and in the context of a study to research the methods and mechanisms currently used to support security and privacy requirements in IIoT platforms another feature model is designed in which such mechanism are assigned to the relevant security and privacy goals. The mechanisms and technologies contained in this feature model partly include the mechanisms and technologies that were already introduced in Figure 20.

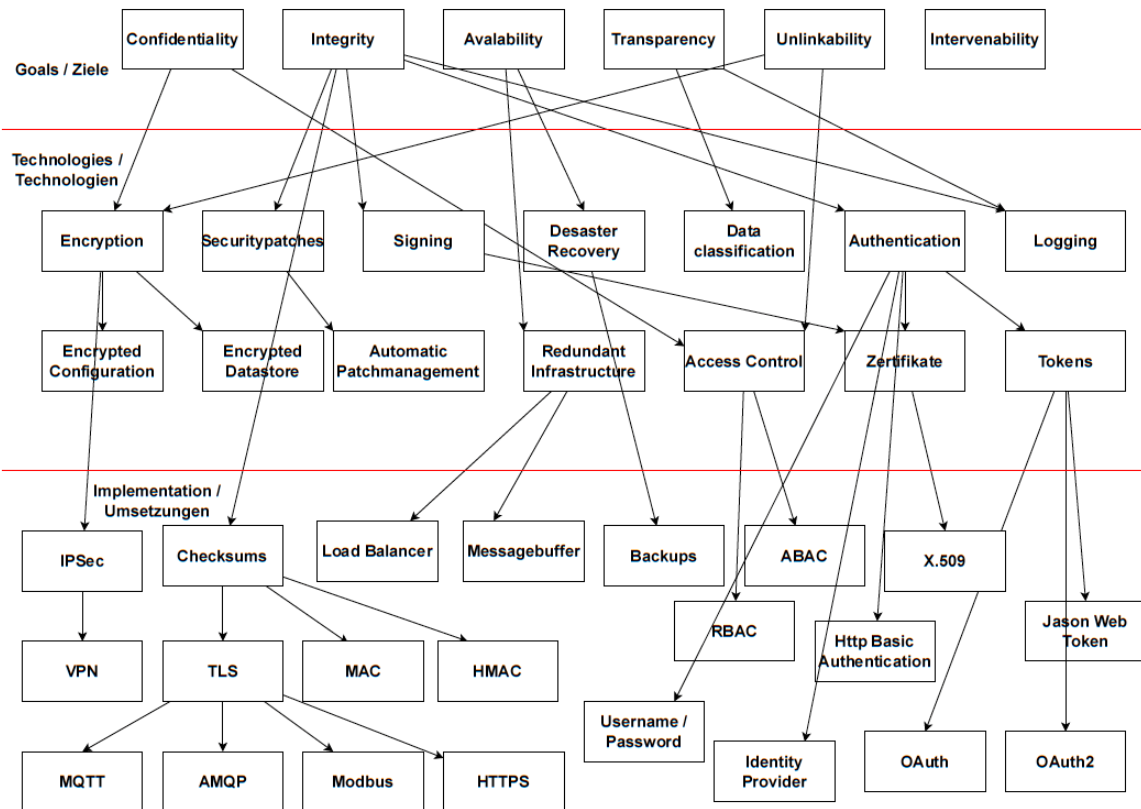


Figure 24: The assignment of mechanisms and technologies (proper to support security and privacy in IIoT platforms) to the security and privacy goals.

In the study to design this feature model, the documentations of several IIoT platform are investigated to obtain the proper mechanisms that potentially can support privacy and security.

Figure 24 shows the goals, the technologies and implementations of security and data protection in form of a feature model. The top level of the model is formed by the protection goals confidentiality, integrity, availability, intervenability, unlinkability, and transparency. Confidentiality, integrity, availability represent central requirements for both security and data protection. Intervenability, unlinkability, and transparency are specific to data protection.

The middle layer of the model represents the technologies used to implement the goals. A technology can be used to fulfill multiple objectives. For example, access control technology is used to implement both confidentiality and unlinkability requirements.

The lowest level of the model consists of concretely named implementations of the technologies of the middle level. The selection of implementations here refers to the explicitly named implementations of the technologies within the IoT platforms studied. The goal of intervenability has not been assigned any technologies or implementations, since the implementation of this goal cannot be realized by an additional component within the platform, but must be considered in the design and architecture of the complete application.

A navigable association from a goal to a technology represents that the technology can be used to fulfill a part of the goal's requirements. Whether a technology is sufficient to fulfill a goal depends on the component under investigation. For example, when implementing the confidentiality goal for data in transit, only encryption technology is needed, not access control. However, for data at rest, both encryption and access control are required to meet the confidentiality goal.

A navigable association from one technology to another technology specifies the use of the technology. For example, there is an association from the technology encryption to the technology encrypted data storage. In this case, the use of the technology is concretized without specifying an exact implementation. In the example given, it is only noted that the data store is encrypted, but not how.

A navigable association from one technology to one implementation represents an explicit implementation of the technology. An association from one technology to multiple implementations indicates alternatives that are not mutually exclusive but serve the same purpose.

For example, the technology encryption can be implemented by the implementation IPsec as well as by the implementation TLS. Both implementations can operate simultaneously and serve the same purpose of encrypting data in transit.

A navigable association from one implementation to another implementation represents that one implementation uses, or builds upon, the other implementation. For example, there is an association from TLS to HTTPS. The HTTPS protocol uses TLS to encrypt connections. The protocols assume that security features are used if possible.

Technologies that were identified supporting the profile shown in Figure 24 are summarized in the Section 7.2.

3.5 Support Layer

The Support Layer aims at providing useful common functions and abstractions to ease the realization of the IIP-Ecosphere platform. Thus, it is more a support library than a full layer, i.e., it does not provide an own AAS representing the interface of the layer. However, even as a support library it is used by the Transport and Connection layer, i.e., the support functionality logically forms an own layer. Below, we detail the AAS abstraction in Section 3.5.1, the network manager in Section 3.5.2 and the lifecycle support in Section 3.5.3. Section 3.5.4 discusses the low level system monitoring support, Section 3.5.5 basic mechanisms for identity management and Section 3.5.6 mechanisms to access program resources on demand. Section 3.5.6 combines the parts and pieces into an AAS creation and usage design pattern that is used in several upstream platform components.

3.5.1 Asset Administration Shell Abstraction

A core aim of the Support Layer is to abstract over the used AAS implementation. This allows for flexibility (the AAS implementation can be exchanged), but also to mitigate risks of impacts by the currently evolving AAS standard and its implementations. Thus, the abstraction described here aims at supporting the application of AAS for the description of interfaces (R7), the application of standards (R2) and enables openness for different AAS implementations, including potential upcoming commercial implementations (R6). Further, an abstraction contributes to the IIP-Ecosphere goal of increasing interoperability, as currently several AAS implementations do exist that potentially do not interoperate (see LNI Testbed Asset Administration Shell⁴⁷). Thus, an abstraction also mitigates development risks, as the current rather dynamic external implementation activities may lead to partially disruptive technical changes.

⁴⁷ <https://lni40.de/lni40-content/uploads/2020/11/AAS-testbed.pdf>

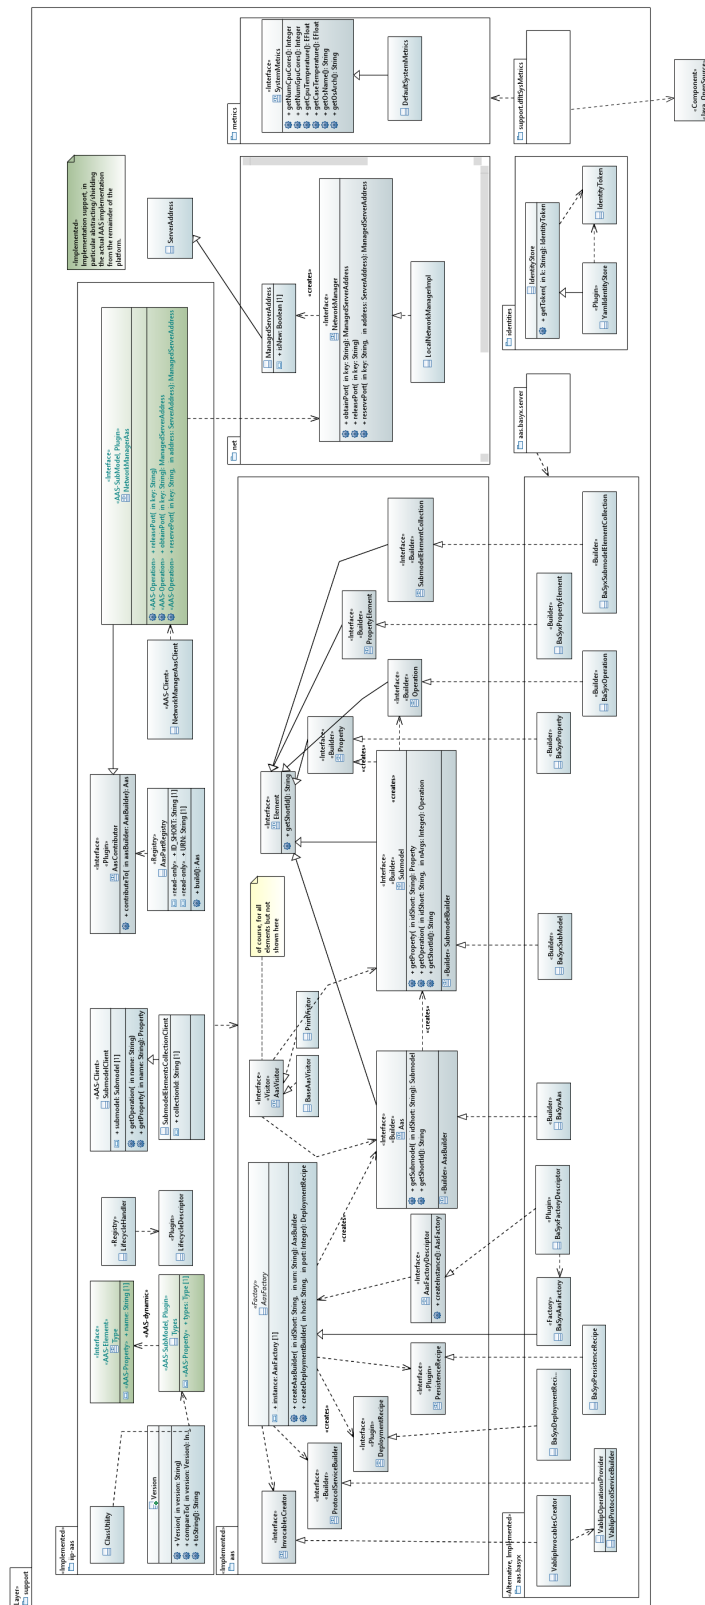


Figure 25: Support Layer overview (only selected classes/interfaces/operations are shown)

Figure 25 depicts the three parts of the support layer. The core is the `aas` component, which defines the IIP-Ecosphere abstraction of AAS. The `iip-aas` component on top utilizes the AAS abstraction to add further functionality that eases the realization of the IIP-Ecosphere platform, e.g., mechanisms how to dynamically link alternative and optional AAS sub-models of different components into the platform AAS. We employ BaSyx as the default AAS implementation of the IIP-Ecosphere platform. The `aas.basysx` component implements the interfaces defined by the `aas` component and provides a

factory implementation to transparently instantiate abstract concepts based on the underlying BaSyx implementation. As BaSyx ships with a large number of dependencies and not all of these dependencies may be needed on an edge device, e.g., when deploying an AAS remotely to a central server (cf. Section 3.1.2) persistent storage to a database is not needed, we aim for a dependency-reduced `aas.basyx` component and an `aas.basyx.server` component including all dependencies.

The `aas` component mainly consists of the instance factory as well as interfaces defining the functionality to be provided by an AAS implementation⁴⁸. It is important to distinguish here between AAS interfaces (such as `Aas`, `SubModel`, `Property` and `Operation` following the AAS meta-model [32]) and the associated (nested) builder interfaces used to build concrete instances of these interfaces. The AAS interfaces provide access to the respective information and, to a certain degree, also allow for modifications, in particular if the interface represents a connected, deployed AAS element. In contrast, the builder interfaces are responsible for creating these instances, allowing for a concise coding style and additional consistency checks, e.g., preventing typical usage errors of the underlying AAS implementation.

Instances of the AAS interfaces can only be created through the factory and the builders, i.e., the top-most AAS-builder can be obtained from the `AasFactory` and all subsequent builders for nested AAS elements (sub-models, element collections, properties, operations) can transitively be obtained from the actual builder. Specific extensions to the typical AAS interfaces are the deployment support (`DeploymentBuilder`), the remote protocol support (`InvocablesCreator` and `ProtocolServiceBuilder`) as well as the `AasVisitor`. The `DeploymentBuilder` aims at realizing and encapsulating typical deployment recipes, such as local or remote AAS deployment. The protocol support encapsulates a specific remote communication protocol to implement the dynamic/active behavior of an AAS. This builder creates function objects delegating the respective operation to the protocol. The function objects can be attached through the `aas` interface to the underlying AAS implementation. This can happen in a straightforward manner if an AAS has direct access to an implementing instance such as a service. However, if AAS and implementation must be separated, e.g., due to remote deployment of the AAS or due to a programming language/process border (e.g., Python for AI services), a (remote) communication protocol must be utilized. As several options for such protocols do exist, e.g., REST, RMI, GRPC, etc., it is not possible just to provide a single protocol abstraction rather than to allow for openness. Therefore, we offer a pair of interfaces, the `InvocablesCreator` being responsible for the function objects to be attached to an AAS (this is just a kind of factory rather than a builder) and a related `ProtocolServiceBuilder` being responsible for building up a server/service instance and registering the actual implementation functions for the function objects. Ultimately, the `AasFactory` is responsible for creating a matching pair of instances for a given protocol.

In addition, the abstraction encompasses an `AASVisitor`. As usual, a `Visitor` allows traversing a data structure in an extensible, polymorphic manner (based on inversion of control) without knowledge about the structure, need for explicit alternatives over types or type casting. Moreover, visitor instances can be applied to any element in the data structure and, thus, perform a partial traversal. Further, there is usually not a single `Visitor` implementation rather than many, each one for a specific purpose. Besides the interface, we provide the `PrintVisitor` which emits the structure of the AAS in textual form in particular for testing/debugging. Further, we provide, as usual, an empty basic implementation, the `BaseAasVisitor` to be used by visitor implementations to handle changes to

⁴⁸ We follow a pragmatic and agile approach here, i.e., we follow the meta-model in [32], but we do not aim to be complete from the very beginning. We add interfaces and operations only on usage demand. Ultimately, at latest at the end of the IIP-Ecosphere project, the abstraction shall be complete with respect to the most recent, implemented version of the AAS specification.

the visitor interfaces in a default manner, i.e., further AAS elements will then not per se lead to a compile error.

Along with the further evolution of the AAS concept, more and more standardized AAS structures will be defined. One such structure is the Technical Data Submodel [2] including manufacturer information, nameplate etc. As an example, the AAS abstraction layer takes up the [2] submodel specification and allows to create and read such structures in terms of specialized submodel and submodel element collection types. This structure is not shown in Figure 25.

A concrete implementation of the AAS abstraction provides an AAS factory. Except for the visitors, which are based on the abstraction rather than a concrete implementation and, thus, can directly be created on purpose by client code, instances of all other concepts can be obtained directly or indirectly from the AASFactory. Concrete AAS factories are supposed to announce/register themselves via the AasFactoryDescriptor and the Java service loader mechanism⁴⁹, so that just the presence of an AAS implementation on the Java classpath enables the abstract AasFactory to create concrete instances.

The default implementation of the AAS abstraction is based on Eclipse BaSyx. The `aas.basyx` component implements the interfaces, typically in terms of adapter/wrapper⁵⁰ classes, i.e., classes that delegate the actual operations to the underlying BaSyx implementation. As remote communication protocol, the default implementation offers an extensible form of the BaSyx Virtual Automation Bus (VAB, in variants TCP, HTTP and HTTPS) through the VabIipInvocablesCreator and the VabIipOperationsProvider. Further, external protocols may be added using the ProtocolCreator (and the related JSL ProtocolDescriptor, both not shown in Figure 25). The default implementation in `aas.basyx` provides a mapping of the Technical Data Submodel [2] to the underlying BaSyx implementation of [2].

In an installation setting, various kinds of AAS servers may be used, e.g., in-memory servers on edge devices or servers with persistent storage of the AAS on a central IT side. However, the different forms of servers imply different dependencies, in particular, database dependencies may not be feasible in resource limited environments such as edge devices as already mentioned above. Thus, implementations of the AAS abstraction are encouraged to reduce dependencies where ever possible to allow for execution in all environments. For IT side installations, all dependencies may have to be included to allow, e.g., for persistent database storage. For this purpose, we separate the AAS implementation into two parts, the (client-side) AAS for all environments and the server side. To announce the server side, we introduce the AasServerRecipeDescriptor (not shown in Figure 25), which, if present, hooks the server component with all its dependencies into the AASFactory and makes such servers transparently available.

The `iip-aas` component paves the way that AAS (sub-models) for the different IIP-Ecosphere platform layers can be collected and deployed as a single representation of a running resource depending on a given deployment mode. Therefore, the `iip-aas` component defines the AasContributor interface and the AasPartRegistry. The AasContributor is a plugin interface supposed to be implemented by upper platform layers to create the respective AAS (sub-model) and to register the implementing function objects with the protocol builders. An AasContributor can indicate whether prerequisites are met so that its AAS can be created. Instances of AasContributor are supposed to be announced/registered via the JSL mechanism. The AasPartRegistry provides access to those plugin instances and, e.g., triggers the creation and the deployment of an entire AAS for an installation. Thus, interfaces marked with the stereotype «AAS» (from the IIP-Ecosphere profile,

⁴⁹ <https://docs.oracle.com/javase/8/docs/api/java/util/ServiceLoader.html>

⁵⁰ https://en.wikipedia.org/wiki/Adapter_pattern

see Section 3.4) are supposed to be implementations of the `AasContributor` interface and to announce themselves via JSL.

As the current specification of AAS and, thus, BaSyx do not contain mechanisms for user-defined types, we will focus on simple property types such as strings in our AAS. Where possible, we also avoid complex types in operation parameters and, if required, use JSON strings to transport multiple values, e.g., objects, arrays or maps. In the future, AAS may allow for modeling such types and BaSyx may provide implicit mechanisms to handle complex objects. Thus, to simplify later code revisions of the IIP-Ecosphere platform and to avoid conflicts with, e.g., annotation-based JSON libraries, we decided to provide some support for JSON marshalling using boilerplate code in the `iip-aas` component, e.g., to handle return values and alternative exceptions for operation calls. Similarly, we did not identify mechanisms to programmatically resolve AAS references in BaSyx, we decided to represent references as Strings carrying the name of an element in a submodel element collection denoted by dependencies or associations or as URLs.

Moreover, `iip-aas` provides common classes to build up parts and pieces of platform AAS instances according to IIP-Ecosphere AAS conventions. Examples are the setup of AAS information via JAML, the resolution of images, e.g., for AAS nameplates, or the `ClassUtility` which turns Java meta-classes into AAS elements and modify the information about available types reflected in the Types sub-model of the IIP-Ecosphere platform AAS. Akin, `iip-aas` implements basic forms of `AAS-client`, i.e., classes that conveniently wrap access to certain AAS parts such as operations or properties. Subclasses shall use or refine the basic functionality to implement concrete accessors, e.g., operation execution.

As far as possible, we aim for a transparent AAS integration. Therefore, platform code **must utilize the abstraction** for the aforementioned reasons. If AAS functionality is not available, new AAS concepts become available or the underlying implementation changes significantly, a revised/extended AAS abstraction may be required, which, in turn, may require changes to existing platform code.

3.5.2 Network Management Support

In addition to the AAS abstraction, the support layer also provides basic network management functionality, in particular for TCP port negotiation. The network manager supports two modes, based on registered and dynamic/free ports. Both modes are relying on a self-selected key for the respective port, e.g., representing a service or a channel/topic identifier. Central services can register themselves with a platform-wide known key. Dynamic services are supported by assigning/reserving free (ephemeral) ports. Furthermore, the network management support can record the number of instances accessing a certain service represented by its known key. This is in particular important if services shall be started/stopped dependent on the actual use, i.e., if no further instance is using a service it can be stopped and the resources can be freed.

Network managers can be stacked, i.e., a parent network manager can contain (more) centrally registered addresses (e.g., for overarching communication brokers) while local managers focus on local (ephemeral) ports. The `NetworkManagerAas` realizes the active AAS frontend network manager instances, in particular for a central platform manager instance. The `NetworkManagerAasClient` implements an AAS-based access to the `NetworkManagerAas`, i.e., to allow implementing components to access a central network manager, also through stacking. It is important to note that not all components rather than installations may require a network manager. Further, not all network managers, in particular not local instances on (edge) resources must be exhibited through an AAS.

BaSyx provides some support for Transport Layer Security (TLS), e.g. HTTPS mentioned above. These mechanisms are available through the support layer and taken up by the platform. However, full use would require key exchanges, e.g., as part of device or identity management, which is still in development. Moreover, the platform provides TLS support for the transport layer (cf. Section 3.6).

For this release, we performed an upgrade of BaSyx from version 1.0.3 used in the last platform release to version 1.3.0. Almost all modifications that were required for this upgrade⁵¹ took place within the support component of the IIP-Ecosphere platform. Moreover, the available test cases for the support layer sufficiently covered the functionality and indicated semantic changes in BaSyx, e.g., how to declare operation parameters, which could be solved by some additional mapping code in the support layer.

3.5.3 Lifecycle Support

A further basic capability is to start up components in a uniform but extensible manner. This is particularly important as individual components may rely on different technology imposing different technological requirements on the startup process. Moreover, it supports the transparent realization of optional and alternative platform components. Therefore, the Support Layer defines the `LifecycleDescriptor` (not shown in Figure 25), allowing components to do the necessary startup/shutdown operations, declare a startup level (priority) and, if required, stop a component. A `LifecycleDescriptor` defines a priority (akin to startup levels in Linux) and may indicate, whether it desires to terminate the execution of the containing platform instance upon a certain event or condition. A `LifecycleDescriptor` announces itself through JSL and is taken up by the `LifecycleHandler`. The `LifecycleHandler` provides generic startup classes for all components, e.g., with or without the ability to terminate the platform instance, which trigger a respective processing of the lifecycle descriptors.

In some cases, components may conflict in their dependencies or behavior. Then, one solution could be to run only parts of a component lifecycle and to execute these parts in an own JVM. Such partial lifecycles can be defined in the platform in terms of `LifecycleProfile`, JLS descriptors that specify a set of `LifecycleDescriptor` instances to be executed when the name of the profile is stated as command line parameter of the component startup. These profiles also allow for virtualization of such partial component lifecycles.

3.5.4 System-level Monitoring Support

System-level properties such as number of CPUs or GPUs, their actual load or temperature are particularly difficult to access in Java. Moreover, edge devices may have vendor specific interfaces including OPC UA or MQTT to access such information. To enable the generic use of such information, also in the platform AAS, we included the required basic access functionality as an interface and a rather simple default implementation into the support layer. Specific implementations can be added via JSL. One example is `support.defaultSysMetrics`, which relies on `JSensors`⁵². One alternative could be `OSHI`⁵³.

The IIP-Ecosphere platform includes an optional system-level monitoring plugin for Phoenix Contact PLCnext, which accesses some system properties like CPU or board/case temperature via GRPC/protobuf provided by PLCnext (starting with firmware released in 2022).

3.5.5 Identity Support

Some mechanisms in the platform require a certain form of authentication, ranging from anonymous over username/password up to X509 tokens⁵⁴, keystores with certificates or (public) cryptographic

⁵¹ Some modifications also affected the Python implementations of the VAB protocol in the service environments (cf. Section 3.7.3). As the support layer (currently) does not contain multi-language implementations of the VAB, this separation is intended and also the modifications there were expected and indicated during test case execution.

⁵² <https://github.com/profesorfalken/jSensors>

⁵³ <https://github.com/oshi/oshi>

⁵⁴ Originally, a generic form of identity tokens was provided by the connectors component, mainly for OPC UA. This now became a more general mechanism of the platform.

keys as well as SSL key managers. However, storing such information in the configuration model or even in code is not acceptable. Therefore, the platform provides an `IdentityStore` with a pluggable implementation. By default, a Yaml file with the identities is read either from the classpath, a file from the home directory of the actual process or a file determined by an environment variable. Moreover, advanced and sophisticated implementations for central identity and authentication token management can be plugged in here. The configuration of components shall refer to the logical name of the authentication token, which is provided (if known) by the identity store. To allow for more flexibility and to ease identity management, several default names, e.g., starting with a specific device name, if not found, the name of a device group, e.g., edges or servers, etc. will be used.

3.5.6 Resource Support

In many cases, programs rely on file resources that must be resolved and loaded at runtime. In Java, this usually happens via the class loader, i.e., Java archive files (Jars) contain such resource files and the Java class loading mechanism provides access to them. However, besides the standard path starting at the root of the archive file, in some cases the packaging of JARs may dictate further paths, e.g., to so called FAT Jars (Java archives in which dependencies are included, partially dissolved into individual files or folders). In the IIP-Ecosphere platform, this can happen when services are packaged into service artifacts. As an unknown number of additional resolution strategies may be required, the support layer provides a class called `ResourceLoader`, which allows registering additional `ResourceResolver` instances directly or via JSL. All platform components are encouraged to utilize the `ResourceLoader` or to contribute required resolution strategies.

3.5.7 Semantic Id Resolution Support

One specific ability of AAS is to mark used elements with a so-called semantic identifier, i.e., a reference to a dictionary detailing what is contained in a certain AAS element. With increasing use of semantic identifiers in the platform AAS, also a resolution of these identifiers becomes important, e.g., on the user interface to display the associated units and descriptions. Of course, within the data processing, also data transformation rules as defined, e.g., by ECLASS⁵⁵ could be used. Besides ECLASS IRDI identifiers, also URL-like IRI are used, e.g., in the specifications of AAS standards. A semantic id resolution mechanism must take care of all those identifiers, potentially taking into account mechanisms implemented by the AAS framework as well as potentially commercial licenses on catalogs and web services as they apply for ECLASS.

The resolution support of the IIP-Ecosphere platform works as follows. A generic semantic id resolution interface provides access to the resolution mechanism. The result of a successful resolution is inspired by the ECLASS dictionary and returns currently the version, the revision, and, in multiple languages, the name, structure name and a free text description of the respective concept. The resolution is done by JSL-based resolution plugins, for which we define two simple standard plugins, both based on a local Yaml-based catalog file: a small excerpt of the ECLASS catalog for the concepts that we use and an excerpt of the AAS IRI definitions from [43]. The IIP-Ecosphere platform also includes an optional resolver performing online resolution through the ECLASS web service relying on the identity management (Section 3.5.5) to access a required authentication certificate. Similarly, a BaSyx-based resolution can easily be realized in the BaSyx component by defining the respective resolver implementation as respective, additional JSL service.

In addition to the two standard plugins based on local catalogues, the IIP-Ecosphere platform also encompasses a semantic resolution plugin based on the ECLASS web service. This web service is commercial and only available with a respective account and security certificates, which are not

⁵⁵ <https://eclass.eu/> IIP-Ecosphere is grateful for the support of Eclass and the ability to use the Eclass catalogue within IIP-Ecosphere in the context of a research license.

included into the source code/binaries of the ECLASS semantic id resolver. This information can be added through the identity support of the platform (cf. Section 3.5.5).

The platform AAS submodel defines an operation to resolve semantic ids, realized by the resolver mechanism discussed above. This allows the UI to request any occurring semantic id and, if resolvable, to use the results to describe individual elements, currently for two languages supported by the built-in catalogs, namely German and English.

3.5.8 Task Tracking Support

For longer running tasks, such as service deployments, tracking and reporting the state of the execution to the user is required, e.g., on a user interface. However, the IIP-Ecosphere platform is a distributed system, i.e., task information must be passed among the executing resources in a manner, that multiple resources can collaborate on the same task. For this purpose, the IIP-Ecosphere platform provides a thread-based task tracking mechanism, which is integrated into the operation execution support for AAS operations.

For modeling AAS operations, we follow the convention, that usual AAS operations are not tracked at all and behave like a synchronous call. Top-level operations that shall be tracked are named with the name suffix “Async”, run in parallel and return their task identification immediately. Lower level operations that can be tracked are marked with the name suffix “ByTask”, offer an additional parameter “taskId” and use the task id for reporting their status. Status reporting is defined as part of the transport layer.

3.5.9 AAS Creation and Usage Pattern

For using the support component in an upstream component or layer, we suggest the creation and usage/design pattern illustrated in Figure 26. As stated in the previous sections, we are using the AAS abstraction (support.aas) as a frontend, i.e., through the AasFactory without direct dependencies to the default implementation of the AAS abstraction (support.aas.basyx) hiding BaSxy (org.eclipse.basyx). The IIP-Ecosphere-specific AAS library (support.aas.iip-aas) provides helpful classes and mechanisms that we use in this pattern.

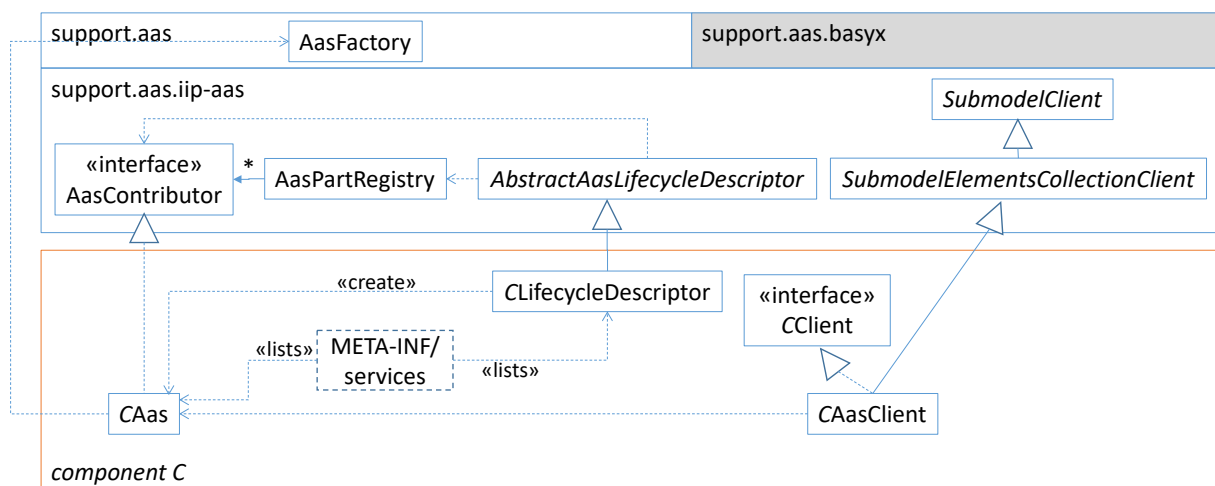


Figure 26: AAS creation and usage pattern involving support layer classes and mechanisms.

To illustrate the pattern, some classes of support.aas.iip-aas as well as some classes of a prototypical component providing an own AAS (*component C*) are depicted in Figure 26. The AAS of *C* is implemented in *CAas* (typically using the name of the component or a suitable shortform as prefix of the class names following Java and platform conventions for the naming, e.g., *C* could be *ecsRuntime*, the AAS could be in *EcsAas*). *CAas* uses the *AasFactory* to create sub-models, properties and operations. However, to be part of the IIP-Ecosphere platform AAS, *CAas* is also an

AasContributor, which defines methods for creating a sub-model (for a given AAS) and for hooking into the AAS implementation server using the recipe interfaces of the IIP-Ecosphere AAS abstraction. To become active, CAas (or the AasContributor, respectively) are mentioned as JSL service in the services directory and, through JSL, become automatically active in the AasPartRegistry (which calls its AasContributor instances appropriately). However, to have a single, central AAS server and to hook the individual parts into that server with the right setup information, we need a lifecycle descriptor. A basic form, that creates also the AAS platform server instances if needed, is provided in terms of the AbstractAasLifecycleDescriptor, which is specialized in component C in terms of the CLifecycleDescriptor. To become active, the CLifecycleDescriptor must be listed as JSL service in the services folder. In this combination, the AAS of component C is created at the right point in time and automatically deployed to or registered with (as stated in the respective setup) the platform AAS. During this creation process, also further AAS may be created, e.g., to represent a device AAS including vendor information [43, 2].

For using the information in the AAS during the execution of other platform components, one could now request the platform AAS instance from the AasPartRegistry and operate on it through the abstraction interfaces provided by support.aas, e.g., to find a certain operation and to call it. However, if all platform parts do that directly, evolving the structure of individual sub-models becomes nearly impossible (or simply a mess). Thus, each component defining a part of the IIP-Ecosphere platform AAS shall also provide a client implementation. For this purpose, support.aas.iip-aas provides two typical basic clients, namely the AbstractSubmodelClient (for properties and operations defined on sub-model level) and the AbstractSubmodelElementsCollectionClient (for an element located in a submodel elements collection in a certain sub-model). The component providing the client shall now define an interface for the respective operations (CClient) and implement that interface in terms of either a specialized AbstractSubmodelClient or an AbstractSubmodelElementsCollectionClient, in Figure 26 shown as CAasClient. Components that want to access the AAS, shall use the client interface and the concrete client implementation. While the CClient interface does not seem to be required here, it helps testing against mocked instances, e.g., in the command interface of the platform.

3.6 Transport and Connection Layer

The Transport and Connection Layer is responsible for connecting resources among each other, with a platform installation on a central IT or even with external cloud environments. We start off summarizing the requirements for the Transport Layer in Section 3.6.1. Then we will turn to the two interrelated components in this layer, the Transport Component (Section 3.6.2) and the Connectors Component (Section 3.6.3). Finally, in Section 3.6.4, we will discuss the realization of the requirements by the two components.

3.6.1 Requirements

In the requirements collection [13], the transport layer is particularly characterized by the requirements summarized in Table 5:

Table 5: Specific requirements from [13] for the Transport and Connection layer (not repeating the general requirements in Table 2 and Table 3).

Requirement	Summary
R13	Connectivity to other actors
R13a	Connectivity with I4.0 devices
R13b	Connectivity with I4.0 platforms
R13c	Connectivity with other IIP-Ecosphere platform instances
R14	Open and flexible connectors
R14a	At least OPC-UA and MQTT connectors

Requirement	Summary
R14b	TCP-IP support
R14c	Bluetooth LE support
R15	Connectors shall be as uniform as possible
R16	Integration of connectors shall be open and flexible
R17	Potential distribution of connectors to various devices
R17b	Management of connectors by platform
R17c	Connectors shall be parameterizable
R18	Securing connectors
R19	Use of a minimum set of internal data formats, examples mentioned in R19a, R19b
R19a	Example input formats (southbound)
R19b	Example input formats (northbound)
R19c	Restful APIs with JSON/XML
R19d	Example output formats (northbound)
R19e	Output data clocked in 5 second intervals
R19f	Data format conversion
R19g	Mechanisms to manipulate data
R20	Application-specific data paths (through the configuration model)
R21	Low impact on data throughput
R22	Platform data throughput of 500 Gbytes per year
R28	Machine pulse of 8 ms
R38	Appropriate authorization mechanism
R40	Role-based access control and TLS
R44	IDS-based connectors (including their security profile)
R49	Appropriate mechanisms to ensure data transfer and data sharing concerning principles of data protection such as legitimated purpose
R66	Pseudonymization and anonymization of data (transferred or shared)

Requirement R19c REST-APIs is not relevant for this layer as service and layer interfaces are expressed through AAS, which in the default implementation form a REST-API. However, JSON and XML mentioned in R19c may be potential wire formats for data transport. Further, as the Transport and Connection Layer shall support (complex) data types in generic manner, the examples mentioned in the explanations of R19a, R19b and R19d in [13] are covered generically. Regarding the general platform requirements in Table 2, in particular, R1, R2, R5, R8 and R10 apply to the Transport and Connection Layer. It is important to note that the realization of (sub-)requirements referring to the configuration model such as R17a or R20 will be discussed in Sections 3.11 and 4, but must be prepared already in the Transport and Connection layer.

With respect to the soft-realtime requirement R10, a stream-based data processing approach seems to be a feasible solution. This impression is supported by the successful application of such approaches in the Big Data domain [28] and the observation that several I4.0 platforms with (soft-)realtime promises typically rely on some form of stream-processing [35]. However, the streaming approach shall not impose unnecessary limitations to the data paths so that, e.g., (AI-)processors can operate with multiple in- and outputs or streams returning to the source (machine connector or underlying platform) can be realized.

3.6.2 Transport Component

At a glance, a Transport Component may appear to be superfluous if we build the platform on the capabilities of the AAS approach. We will outline our rationales first and then turn to the design of the Transport Component.

Initial experiments [39] with Docker containers and AAS (using BaSyx version as of July 2020⁵⁶) on Raspberry Pi 3⁵⁷ as well as on Phoenix Contact PLCnext edge devices showed that the typical response time of operations without computational load is around 23 ms. In contrast, property value accesses can be executed at 4-10 ms. For comparison, plain Java Remote Method Invocations operate in this setup at 2-4 ms. Further investigations of this discrepancy indicated that utilized components from the BaSyx examples represent one operation as three internal REST calls, where two of them are required for managing internal functionality. Using lower level BaSyx components, e.g., through the `InvocablesCreator` and `ProtocolServiceBuilder` discussed in Section 3.5, operation calls can be executed at the same response time as property accesses, i.e., around 4-10 ms. This seems to be promising for R10 and, in particular, the 8 ms machine pulse mentioned in R28. However, the measured times are not suitable for (soft-)realtime processing or software-based stream-processing.

Consequently, we will use AAS primarily as control interface for the infrastructure layers and the platform services (R7). For cross-linking services in terms of data streams we will rely on a dedicated streaming approach, as to our very knowledge so far no concepts are provided for streamed data among multiple AAS⁵⁸. In such a streaming setting, the integration of services (potentially implemented in different programming languages, R113a) could be done via AAS (not preferable as argued above) or through a specific streaming data interface. In the latter case, the service AAS is used to describe the streamed data types and the data connectors³⁸ but not for the actual streaming transport. While AAS may be preferable for uniformity (R7), specific streaming interfaces will allow for better performance (as we will detail in Section 3.6.2.3).

We structure this section into a brief review of related streaming approaches leading to some technology candidates, the design of the component and its (initial) technical validation.

3.6.2.1 *Related Approaches*

We discuss now approaches in the field of stream-processing, their relation to I4.0 and whether they could be useful for realizing the IIP-Ecosphere platform. As often certain protocols are required in I4.0 settings, we further discuss protocol realization candidates.

Regarding stream processing, we briefly review now related approaches, in particular stream processors (with own resource management approach) as well as stream processing libraries (focusing more on the stream-based transport). It is important to make this distinction, as the dynamic deployment and the adaptation capabilities foreseen for the IIP-Ecosphere platform shall integrate with rather than contradict automated management capabilities of the selected stream processing approach. From previous experience we know, that modifying a stream processor such as Apache Storm to introduce desired adaptation capabilities can be a tedious work that, if not accepted by the developing company, will not scale/evolve with future developments of the stream processor. Besides technical capabilities, important aspects are R13 (utilized transport protocols), R14 (openness and flexibility, in particular for connectors, transport protocols and formats) as well as availability and licenses (R5) of the individual approaches.

⁵⁶ At that time there were no version-based releases of BaSyx.

⁵⁷ Raspberry Pi is frequently used as IoT mockup device in literature and practice, e.g. for cost-effective showcases or even in experimental evaluations.

⁵⁸ AAS are designed to describe production assets along their lifetime rather than software components or even soft-realtime data services. Thus, this statement shall not be considered as critics rather than a potential limitation of AAS that has to be mitigated in IIP-Ecosphere by a different solution or additional technical means.

Table 6: Selected scientific streaming processing approaches (related to IoT)

Name	Mgt.	Based on	Protocols	Edge	Availability	License	Latest
EdgeWise [16]	Yes	Apache Storm		Pi3	Yes	APL2.0	2016
Frontier [29]	Yes		“WiFi”	Pi	-	?	?
[7]	Yes	Apache NiFi		Pi	-	?	?
Dart [6]	Yes		REST, JSON	Pi	-	?	?
PESP [20]	Yes	RabbitMQ			-	?	?
VISP [19]	Yes	Spring cloud stack, RabbitMQ	AMQP, MQTT		Yes	APL2.0	2018
Esc [34]	Yes				-	?	?

Table 6 compares related approaches that can be found in literature on stream processing for edge devices or IoT. Three of the approaches were evaluated on Raspberry Pi devices (mocking IoT devices, typically in rather different IoT application settings). For two of the approaches the implementation is available, and only [19] utilizes IoT protocols (as provided by the underlying technology). All of the approaches provide more or less (automated) management/deployment functionality. Although potentially interesting in an evaluation context, we see existing management functionality as a potential risk as they usually are designed for a certain setting and integrating, interfacing or in the extreme case replacing such functionality may lead to unforeseen complications.

Table 7 summarizes a set of recent stream processing frameworks and libraries that were identified through a web search (without specific focus on IoT/edge). Not included are commercial approaches like Grovestreams, Hazelcast, or Amazon Kinesis that were also part of the result set as R1, R5 and R6 in [13] express a clear direction towards Open Source.

Table 7 consists of two parts, the upper part for IoT related approaches and the lower part for generic stream processing approaches. In these approaches, the presence of management functionality correlates with the nature of a stream processor rather than a streaming library. Moreover, stream processors are usually realized as a complete stack, usually based on a cluster-based installation, e.g. with centralized servers such as Apache Zookeeper or a resource manager. Within the IT server installation of an IoT platform, such setups may appear to be feasible. However, we expect serious limitations for devices close to the production / OT⁵⁹ side. It is worth mentioning that some sources indicate that approaches like Apache Flink try to remove resource consuming central services like Apache Zookeeper and even consider the realization of specific versions for edge devices (e.g., a 2021 release of Flink shall include edge support). From a research perspective, e.g., in [16], such widely used stream processors are also criticized for their inflexible adherence to the One-Worker-Per-Operation-Architecture (OWPOA) as this design loses efficient processing scheduling opportunities by relying on the underlying operating system scheduler. A positive trend seems to be that more and more security features are built into the frameworks. Examples of such security features and a potential methodology to integrate them in a framework—especially from the early phases of framework design—is introduced in Section 3.9. This is in contrast to the initial tendency of, e.g. Apache Storm, to sacrifice security and encryption for throughput. Except for approaches dedicated to the IoT domain, full-stack frameworks typically realize own (usually fixed) transport protocols and, if at all, realize own

⁵⁹ As discussed in [13], **OT** refers to **O**peration **T**echnology, i.e., the control and monitoring of production machines, which typically operate under hard-realtime conditions. In contrast, **IT** (**I**nformation **T**echnology) such as the IIP-Ecosphere platform typically can only operate under soft-realtime constraints. Nowadays, edge devices may bridge OT and IT, e.g., in terms of separated, but integrated hard- and soft-realtime cores, potentially controlled by different operating systems/software.

internal interfacing concepts, e.g., based on REST-APIs. Thus, as far as we can see, (flexible, exchangeable) IoT-based transport protocols are rather uncommon and, typically, AAS is not considered for interfaces at all.

Table 7: Selected stream processing frameworks and libraries (*=incubating, vispl = visual programming language)

Name	Mgt.	Language	Protocols	Env	License	Latest	Comment
Apache Edgent	Yes	Java	MQTT, ...	edge	APL2.0	2017*	Own language
Apache Streampipes	No	Java, Typescript	MQTT, OPCUA, ROS, ...	IoT	APL2.0	2020*	ML support, Kafka, container, vispl
Benthos	No	Go	AMQP, MQTT		MIT	2020	Own language
Eclipse IoT Streamsheets	Yes	JavaScript	MQTT, OPCUA, ...	IoT	EPL2.0	2020	Spreadsheets, Docker
Eclipse Kura	Yes	Java	MQTT, OPC-UA, ...	IoT/edge	EPL1.0	2020	OSGi, Docker, vispl, stack
EdgeX Foundry	Yes	Go, C	MQTT, OPC-UA, ...	IoT/edge	APL2.0	2020	REST, Stack
Flogo	Yes	Golang		IoT/edge, Cluster	BSD3-Clause	2019	TensorFlow, zero code
Sensorbee	No	Go		IoT	MIT	2017	ML integration, own language, documentation
Akka	Yes	Scala, Java	HTTP	Cluster	APL2.0	2020	Actors
Apache Apex	Yes	Java	Hadoop RPC	Cluster	APL2.0	2018	YARN
Apache Beam	Yes	Java, Python, Go		Flink, etc. cluster	APL2.0	2020	Requires processing cluster
Apache Flink	Yes	Java, Scala		Cluster	APL2.0	2020	Zookeeper, edge 2021?
Apache Flume	Yes	Java	Avro, protobuf	Cluster	APL2.0	2019	Zookeeper, big footprint
Apache Gearpump	Yes	Scala			APL2.0	2019*	Storm/Samoa compatible, YARN
Apache Kafka	Yes	Scala, Java		Cluster	APL2.0	2020	Zookeeper
Apache Kafka Streams	Yes	Java		Cluster	APL2.0	2020	Requires Kafka
Apache NiFi	Yes	Java		Cluster	APL2.0	2020	FlowFiles, REST, vispl
Apache Pulsar	Yes	Java		Cluster	APL2.0	2020	Brokers, Bookkeeper, Zookeeper
Apache Samza	Yes	Scala, Java		Cluster	APL2.0	2019	YARN, Kafka

Name	Mgt.	Language	Protocols	Env	License	Latest	Comment
Apache Spark	Yes	Scala	Spark RPC	Cluster	APL2.0	2020	shared memory
Apache Storm	Yes	Clojure, Java		Cluster	APL2.0	2020	Zookeeper
Hortonworks Streamline	No	Java, JavaScript	REST		APL2.0	2018	Documentation
Spring Cloud Stream	(No)	Java	open*		APL2.0	2019	Depends on Spring
StreamFlow	Yes	Java		Kafka, Storm cluster	APL2.0	2015	Vispl for Kafka, Storm
Streamtz	No	Python			Continuum	2020	Pandas, cuDF

Some of the approaches (Apache Streampipes, Sensorbee, Streamtz, Flogo) listed in Table 7 provide Machine Learning (ML) support/integration while other approaches already realize concepts for the integration of programs in multiple languages, e.g., Apache Storm and Apache Spark for Python. Approaches like Apache Streampipes, Eclipse Kura, Flogo, Apache NiFi, or Streamflow take up the trend towards visual programming, low code or even no code (as also identified for some I4.0 platforms in [35]). Although convenient, an integration of configuration modeling with existing visual programming languages for the orchestration of services may be a risky approach.

Summarizing these findings, only few approaches remain as candidates for the realization of streaming in the Transport Layer, which may involve edge devices, servers and even clouds. This (as well as license and normative considerations) requires a flexible selection of the transport protocol (R14) and the wire format. Moreover, the requirement of using AAS wherever possible (R7) or at least to allow for an exchange of the communication protocol excludes almost all stream processing frameworks (of course not from comparative evaluation settings). Needless to say that a candidate approach shall be able to handle synchronous (one input data item is related to one or no output item) and asynchronous (inputs and outputs can be decoupled) processing of data items as well as resources local (among local processes) and external network communication, e.g., for input and output.

As a result of these filter criteria, Apache Streampipes, Sensorbee, Hortonworks Streamline or Spring Cloud Stream appear to be feasible candidates. However, initial experiments indicated serious problems with Sensorbee and Hortonworks Streamline ranging from incomprehensible or non-existing documentation to problems when executing the respective examples. Apache Streampipes (APL2.0) ships with interesting features, several kinds of generic connectors, exchangeable transport protocol and wiring format although it is still considered to be in incubation state. As alternative we see Spring Cloud Stream⁶⁰ (APL 2.0), which allows exchanging the transport protocols for individual in/out-bound streams, supports user-defined payload wire formats, flexible exchange of communication protocols, network properties per data path among processors, implicit payload conversion (also through our serializer mechanism) and dynamic stream re-routing at runtime. Moreover, Spring Cloud Stream was successfully applied in [19].

⁶⁰ <https://spring.io/projects/spring-cloud-stream>



Figure 27: Transport Component overview (comments partially cropped)

However, it is important that we want to integrate the streaming approach with the connectors and the ML processors in a model-based manner similar to our work for Apache Storm in [9], here with an even stricter focus on isolating the utilized streaming approach. We believe that relying on glue code generation allows us to replace (within limits, e.g., always assuming a data flow graph with some kind of source, processor and sink) the stream approach. Thus, we see Spring Cloud Stream as a good and justified initial candidate for the reasons stated above as well as because it ships with several transport protocols including protocols for public clouds (although it also relies on a large dependency tree particularly induced by the Spring Framework⁶¹). In later stages of the project, we may take Apache Streampipes or an edge-enabled version of Apache Flink into account.

Regarding IoT protocols, several implementations are available, in particular from different projects of the Eclipse IoT⁶² ecosystem (provided under compatible licenses for IIP-Ecosphere). While some projects focus on specific protocols, e.g., Eclipse Paho⁶³ on MQTT, others already integrate various protocols such as Eclipse Hono⁶⁴. Although such integrations may be an interesting foundation, they often rely on specific assumptions, e.g., Eclipse Hono collects binary payload from different protocols and forwards the payload to a fixed default protocol (MQTT). While such approaches may provide access to different low-level protocols or machine connectors (cf. Section 3.7), they may also introduce limitations due to their design choices or do not provide mechanisms for turning such generic implementations into application specific solutions, e.g., through application-specific data translators. Moreover, some transport protocols are currently not applicable, e.g., we currently do not consider OPC UA PubSub due to a lack of feasible implementations, where Eclipse Milo⁶⁵ currently does not support the required OPC UA version.

3.6.2.2 Design

Figure 27 depicts an overview of the packages and (top-level) classes in the Transport component. The Transport component is intended to be deployable as re-usable component rather than to act as a standalone communication container. The main concepts in this layer are:

- The `TransportConnector` allowing to bind transport protocols into the infrastructure. A transport connector allows sending/receiving of data on (virtual) channels. As receiving usually happens in asynchronous manner, implementations that rely on a `TransportConnector` are informed via the `ReceptionCallback` about received data.
- The actual wire format to be used for transport may differ from protocol to protocol. For example, low level transport protocols such as MQTT or AMQP support arbitrary binary payloads (might be with individual size restrictions) while higher level protocols such as OPC pub/sub define their own payload format. However, to be open and flexible with respect to the wire format and to utilize a minimum of data formats within the platform (R19), we foresee a mechanism for data transcoding. For performance reasons, data transcoding shall happen only when actually needed. Specifically, for binary wire formats, the `Serializer` transcodes programming language objects into a binary representation and back. More generically, a `Serializer` is a `TypeTranslator` that can be applied also in other situations, e.g., data processing. In turn, `TypeTranslator` is a combination of `InputTypeTranslator` and

⁶¹ Native executables are in experimental development and may help optimizing the deployment/performance: <https://www.heise.de/news/Java-Framework-Native-Spring-Anwendungen-laufen-ohne-die-JVM-5078681.html>

⁶² <https://iot.eclipse.org/>

⁶³ <https://projects.eclipse.org/projects/iot.paho>

⁶⁴ <https://projects.eclipse.org/projects/iot.hono>

⁶⁵ <https://projects.eclipse.org/projects/iot.milo>

OutputTypeTranslator with cross-over template bindings⁶⁶. Intentionally, we leave the actual technical approaches for transcoding open here (some candidates are JSON, OPC-JSON or protobuf⁶⁷). The actual instances depend on the data types used in the application and are supposed to be generated from the configuration model. While instances of TypeTranslator are supposed to be attached where needed (and may be combined with Serializer instances), Serializer instances shall be usable dynamically on-demand, e.g., for a certain TransportConnector implementation. For this purpose, we provide a SerializerRegistry.

- The TransportConnector instances shall be available to other components of the platform where an internal data protocol is needed. To obtain TransportConnector instances, we define a TransportFactory and exhibit the actual protocol, the wire format and the broker data connector(s) from the platform configuration in the Transport AAS.
- Three default protocol plugins are shipped with the IIP-Ecosphere platform, namely MQTT v3 (based on Eclipse Paho), MQTT v5 (also Eclipse Paho) as well as AMQP (based on the RabbitMQ AMQP client). Each protocol plugin is an own alternative component, the installed ones determine the TransportFactory behavior through a JLS descriptor. The default protocol plugins support optional Transport Level Security (TLS) and, thus, contribute to the realization of R40.
- The streaming approach is already relevant to the Transport Layer as transport protocols and wire formats must be provided accordingly. However, as discussed above, the streaming approach shall also remain exchangeable through glue code generation. Thus, the platform provides also transport plugins for the default streaming approach (Spring Cloud Stream), the so-called Binders, which are realized in turn through the Transport Component. A basic spring component implements convenient mechanisms for applying Spring Cloud Stream in IIP-Ecosphere, e.g., to add serializers to the SerializerFactory through the component setup (in Spring application.yml, to be instantiated from the platform configuration) or to bind the SeralizerFactory to the data conversion mechanism of Spring Cloud Stream (SerializerMessageConverter). In addition, Spring Cloud Stream ships with generic serialization approaches, e.g., for JSON or XML that may be used out-of-the-box. By default, the IIP-Ecosphere platform ships with five alternative Spring Cloud Stream protocol binders for MQTT v3 (based on Eclipse Paho and HiveMQ-client), MQTT v5 (based on Eclipse Paho and HiveMQ-client) and AMQP (based on the RabbitMQ AMQP client). These binders support optional Transport Level Security (TLS). Alternative binders for the same protocol are mainly provided for performance comparison.
- The transport component defines several global platform streams (StreamNames), e.g., for status (StatusMessage), alert (Alert) and trace (TraceRecord) messages or, as forward declarations, for upstream components (to avoid unnecessary dependencies, cf. Table 25 in Section 8.1). The status notification mechanism informs interested parties when containers or services are dynamically added or removed. The notifications consist of a message data structure, which is sent on a pre-defined transport channel. Further, status notifications may report on their progress and may be task-related, i.e., carry a task id (cf. Section 3.5.8). As task-based execution does not have a result, status messages may carry result or failure information in this case. Alerts are created by monitoring components to signal abnormal or undesired

⁶⁶ At a glance, TypeTranslator shall be sufficient, but in some situations, it is convenient that only the required direction must be implemented rather than both. This is in particular true for the machine/platform connectors, which require either direction for different types but usually not both directions. As TypeTranslator inherits from the input/output type translators, it is also possible to use a fully-fledged TypeTranslator in these situations.

⁶⁷ <https://developers.google.com/protocol-buffers>

situations. Traces make the operations of the platform visible. Moreover, the transport component defines a global instance of the default `TransportConnector` and send methods, that may queue messages until the transport connector can be utilized.

It is important to mention that further protocol binders are available for Spring Cloud Stream, e.g., for RabbitMQ, Amazon Kinesis, Google PubSub, Solace PubSub, Azure Events Hub, Apache RocketMQ. These binders may be helpful for realizing Cloud integrations, e.g., in the Storage, Security and Data Protection Layer. However, for uniform usage of protocols within the platform, a respective `TransportConnector` shall be provided (the AMQP connector may already be used for RabbitMQ). Furthermore, it is important to mention that we do not prescribe the amount or deployment strategy for communication servers (Brokers for the mentioned concrete protocols) within a platform installation. The platform configuration shall provide opportunities to define multiple brokers (to be reflected in the Transport AAS) while the broker(s) to be used shall be instantiated through the platform configuration or the network managers into the respective deployment units. Moreover, based on the provided mechanisms of the protocol implementations and the streaming library, different levels of resilience or recovery can be realized, while failover to alternative broker servers may require additional implementation work.

It is also notable that Spring Cloud Stream (with exposed web management interface) and Eclipse BaSyx do collide on the use of Tomcat when being executed in the same JVM. Thus, also for this reason, services shall be executed in their own JVMs and run at least in different processes than other IIP-Ecosphere components to be introduced later, e.g., the service manager, or the ECS runtime.

3.6.2.3 Validation and Evaluation

We discuss now briefly the validation of the design and the implementation of the Transport Component as it has a major impact on the performance of the entire platform. We start off with a discussion of the regression testing approach and turn then to an initial performance evaluation.

The implementation of the Transport Component is subject to **regression testing** and continuous integration. Testing protocol integrations requires some form of server or broker instance. Therefore, further Open Source components are utilized so that the tests are self-contained, e.g., embeddable protocol brokers to simulate the platform side in the respective tests. The required dependencies are only active in testing, i.e., they are not part of a platform installation and, thus, here relaxed license or Java version rules may apply if needed. In the regression tests, we use protobuf and a simple JSON implementation for serialization as well as Apache HiveMQ or Moquette as MQTT broker and the Apache Qpid broker as AMQP broker.

For the Spring Cloud Stream binders we realized a simple setup validating the envisioned streaming capabilities mentioned in Section 3.6.2.1. This is reflected in the communication setup shown in Figure 28. Ingested by a Source (the regression test), a mocked stream component (Transformer) modifies the data (synchronously) and passes the data to the broker (representing the IIP-Ecosphere platform/server). The communication between these instances is handled by the Protocol Binder under test as well as the Serializer selected by the test. The Protocol Binder is based on the respective protocol implementation and in the test bound against a corresponding embedded test server/broker. To test also the flow back, a shortcut client based on a corresponding `TransportConnector` receives the data and ingests modified data asynchronously, which now flows through the Broker, the Serializer and the Protocol Binder back to the Source acting also as Receiver. Combining Source and Receiver is a relevant setup, as a machine/platform connector (to be discussed in Section 3.6.3) also ingests data and may receive information, e.g., to reconfigure an edge device or a machine. The regression test has access to the sent/received information and, thus, can validate the entire flow.




Figure 28: Regression testing data flow for the Transport Component.

In addition, it is also important to understand the (early) fulfillment of **quality requirements**. We determine the respective properties in terms of a **performance experiment**. Figure 29 details the setup of this experiment, which in fact is a variant of the regression test setup. Here, the Source produces a stream of data items at a certain ingestion frequency. Each data item consists of at least 50 values with repeatable characteristics (R19a). We concentrate on the payload and scope out meta-information (R79) for now. A simple Anonymizer takes a produced data item and turns one property (a name String) into simple pseudonyms. An “AI-Service” inspects the data and sends for 5 received data items one “command” back to the Source. Again, on the forward flow, the processors operate synchronously, while the backward “command” flow is ingested asynchronously. The number of received data items is recorded in all processors by simple monitoring probes and written in parallel once per second to a log file. An additional stream is used to asynchronously send experiment control commands to all involved processors, e.g., to terminate the experiment and to close the monitoring log. Items on the experiment control stream are not recorded by the probes.

The processors in Figure 29 can be executed locally (in one process, in multiple processes) or distributed on separated hosts as indicated in Figure 29. For the distributed execution, two brokers are used, one in the local realm and a remote broker in the platform realm. In the local realm, we currently use the same transport protocol/mechanism as in the platform realm, i.e., we focus at the moment on an Inter-Process Communication (IPC) setup rather than an edge setup where at least one stream goes to a different resource or the platform. Replacing the transport protocol, using different brokers or exchanging the wire format for serialization may be subject to future experiments. In this experiment we focus on the basic transport characteristics of the utilized approach.




Figure 29: Performance testing data flow for the Transport Component.

For executing the experiment, we use a selection of the binders available in the platform (HiveMq v3, v5 with QoS AT_LEAST_ONCE, AMQP) with the setup as shown in Figure 29 and a respective (local, embedded) broker (Apache HiveMQ 2020.4, Apache Qpid 8.0.2). As baseline, we realized a plain network communication binder/distributed broker based on Netty⁶⁸, an asynchronous networking library, and the network port management of the platform. For the source, we use a message ingestion rate⁶⁹ per experiment and vary it from slow pace (R28) up to congestion. As wire format, we use a

⁶⁸ <https://netty.io/>

⁶⁹ The ingestion is based on the Spring Default Poller, which is controlled by a fixed delay between message ingestion time slots (translates to a minimum ingestion rate) and a maximum number of messages ingested within a slot (determines a maximum ingestion rate). The effective ingestion rate is within the minimum and maximum ingestion, but subject to an internal congestion control of Spring Cloud Stream.

simple JSON serialization (leading to 650 Bytes of payload). We run the experiment for 1 minute and exclude by default the first three seconds as well as the last second where fluctuations due to network, just-in-time compilation and broker startup activities may occur. Further, some time may elapse until the average throughput is established, which we consider in this experiment as part of the stable measurements although it may significantly cause variations.

The measurements for this initial experiment have been taken on an Intel Core 7-8750U @ 1.90GHz with 32 Gbyte running Windows 10 and OpenJDK 13+33. As we aim at the moment for initial measures, we do not pay specific attention to a clean setup, e.g., getting rid of potentially other process influences such as a virus scanner or system updates.

Figure 30 illustrates the average ingestion rate at the source on the horizontal axis and the average arrival rate at the sink on the vertical axis. Until an ingestion rate of around 1000 messages per second, all binders scale similarly. Over 1000 messages per second, the behavior of the four binders differ significantly. The arrival rate of the MQTT v3 binder starts dissociating from the ingestion rate at around 1500 messages per second. For MQTT v5 this happens at around 2100 messages per second and for AMQP at a rate of roughly 2300 messages per second. While the MQTT v3 binder tries to cope with the ingestion rate until 6500 messages per second (dropping at the sink to 1400 messages per second), the MQTT v3 and the AMQP binders stop operating around 2700 messages per second. In contrast, the experimental Netty binder scales well until 7200 messages per second. Then the sink rate starts dissociating from the ingestion rate and above 9300 messages per second the simple experimental broker implementation stops operating as indicated by the trendline in Figure 30. Moreover, there are noticeable differences in settling time for the average throughput (not shown in Figure 30): All binders require more than 10 seconds to reach the respective average throughput, while Netty requires higher settling times for lower ingestion rates and AMQP leads faster to a stable throughput than both MQTT versions.




Figure 30: Average stream throughput measures for the four utilized alternative binders with trend lines.

As Figure 30 relates source and sink throughput rates, it does not reflect the total number of translated messages. Due to the streaming setup, the messages among source, processors and sink and also messages on the “command” channel (one item per five input messages) are communicated. Thus, the absolute number of transmitted messages per second is higher (least around factor 3.2). Table 8 details these numbers for the measured protocol-client-server combinations. In particular, our HiveMq

readings amount to similar ranges as reported in [25], where two server machines with up to 16 CPU cores but no stream processing approach were used.

Table 8: Total number of translated messages per second in best source/sink transmission situation.

Total number of translated messages per second	
MQTT v3: HiveMq, HiveMq embedded server	6172
MQTT v5: HiveMq, HiveMq embedded server	8908
AMQP: Rabbit MQ client, Qpid embedded server	9531
NETTY	30298

In summary, the required rate of 125 messages at 8 ms machine pace (R28) is supported by all brokers and works in combination with the Spring Cloud streaming approach. At around 50 values per message (650 Bytes of payload in a JSON serialization), a stable ingestion of 1000 messages per second leads to (calculated) 2.1 GByte of data transmission per hour. Moreover, the Netty binder can cope with (calculated) 15.6 GByte of data, which even qualifies for R91⁷⁰. It is important to emphasize that we focus here on pure IPC transport characteristics without significant data processing load. Moreover, we use a single stream, i.e., multiple (moderate) input streams from different edge devices may easily aggregate to even higher frequencies and volumes. In a realistic setting, we expect a multi-server setup as platform installation and potentially also a redundant cluster-based message handling for individual tasks, e.g., in the data integration, so that the envisioned approach qualifies for the given data (transport) quality requirements, in particular frequency and volume.

Further experiments indicate that the discussed behavior is similar when running the data processing within a single JVM, i.e., as threads, or in separate processes. Measurements on real edge devices with inter-device (cross-realm) network communications are subject to future work. As soon as further parts of the platform are available that potentially impact the data size or the performance (meta-information, security, etc.), further experiments shall be performed.

3.6.3 Connectors Component

The Connectors Component is responsible for the communication with already installed platforms (the virtual platform aspect) or machines (potentially connected via some form of edge devices). The aim here is to allow for a bi-directional, typed communication represented in terms of connector instances. Relying on the design of the Transport Component, it is desirable that the machine/platform connectors utilize type translators or serializers for the inbound communication, i.e., to translate received information (if feasible already filtered in application-specific manner) into application-specific datatypes that can further be processed in the IIP-Ecosphere platform. For the outbound direction, (AI-)services or humans may make decisions about changes in the connected machines/platforms. These decisions are represented as information, e.g., commands, and are translated/sent through the connector to the machine or platform. Here, type translators shall turn the application-specific data types received from the platform side into information suitable for the external side. As stated in Section 3.6.2, application-specific type translators shall be realized by code generation to ease the development of applications.

The connectors discussed in this section may be utilized within the realm of the same factory, i.e., they may run at reduced or even no security measures. The connectors may also link to external realms, e.g., via the internet. In this case, adequate encryption mechanisms may apply or even the machine/platform connectors may have to be extended through IDS functionality.

⁷⁰ Based on the transferred messages in Table 8, this leads to 13.5 GBytes up to 66 GBytes per hour.

3.6.3.1 Related approaches

Regarding IoT protocols, several implementations are available, in particular from various projects of the Eclipse IoT ecosystem. Some projects focus on specific protocols, e.g., among others Eclipse Paho on MQTT, Eclipse Milo on OPC UA, Eclipse Californium⁷¹ on CoAP, Eclipse Leshan⁷² on LWM2M, or Eclipse Tahu⁷³ on legacy SCADA/DCS/ICS protocols. Other projects already integrate various protocols such as Eclipse Hono, Eclipse Agail⁷⁴, Eclipse Kapua⁷⁵ with a cloud focus based on MQTT transport or Eclipse Ponte⁷⁶ for mapping IoT protocols into REST. Although such integrations may be an interesting basis for our work, they already realize concepts and ideas that do not fully comply with the requirements of the IIP-Ecosphere platform. For example, Eclipse Hono collects binary payload from different protocols and forwards the payload to a default protocol (MQTT) without options for filtering/translating the payload or for supporting alternative protocols. Similarly, Eclipse Ponte focuses on REST for internal communication, which from our point of view is just one potential alternative. Further, Eclipse Agail emphasizes the cloud aspect but neglects local resources or edge devices. While these and similar approaches may ease the access to different low-level protocols or machine connectors (cf. Section 3.7), they usually do not provide mechanisms to enable core IIP-Ecosphere functionalities such as data filtering or translation. Moreover, if they implement a stack or multiple integrated layers, they usually do not offer AAS functionality (R7).

While it makes sense to review these approaches to find a feasible abstraction as well as to consider existing abstractions and protocol implementations as those mentioned above, it is not productive just to implement connectors to achieve a high number of protocols (this is one of the strengths of existing I4.0 platforms as shown in [35]). It is more important to develop and evaluate concepts to enable openness and extensibility for inbound/outbound directions. From a resource perspective, it is important to realize connector types for the actual needs of the stakeholders (R14a states MQTT and OPC UA). Moreover, one goal is to demonstrate how model-based generation can turn such generic connectors into efficient and application-specific mechanisms already at the bottommost layer of a platform.

3.6.3.2 Design

For the design of this component, it is important to recall that in contrast to the Transport Component, the Connectors Component already deals with processing and translating application-specific data. For example, it is not performant to just ingest, e.g., an entire OPC UA namespace upon each data modification or, if polling/sampling shall be applied, in each poll cycle. It is more important to select the required data in an application-specific manner and to focus on the information that is required by an application running on the platform. We call the step of translating an outbound protocol into an internal protocol (and back) “protocol adaptation”, i.e., a (generated) plug-in ProtocolAdapter will be responsible for this task. One form of implementing the protocol adaptation is in terms of existing TypeTranslator and Serializer instances from the Transport Component, either as the realizations are part of the platform and can be re-used or because they are defined as part of the application and can be generated or are provided as hand-crafted components. However, also other forms of type translation may occur. This applies to connectors that handle generic payload (where the payload format must be translated to application-specific instances and can optionally be filtered/translated). Further, it applies to connectors that are based on a specific information model, such as OPC UA or AAS. In the latter case, we aim for specific TypeTranslator instances that are

⁷¹ <https://projects.eclipse.org/projects/iot.californium>

⁷² <https://projects.eclipse.org/projects/iot.leshan>

⁷³ <https://projects.eclipse.org/projects/iot.tahu>

⁷⁴ <https://projects.eclipse.org/projects/iot.agail>

⁷⁵ <https://www.eclipse.org/kapua/>

⁷⁶ <https://projects.eclipse.org/projects/iot.ponte>

linked to a generic model interface abstracting over the underlying information model. However, not all approaches support the same range of concepts and types, e.g., OPC UA allows different kinds of custom datatypes while AAS does not. Thus, connectors will differ in the offered functionality of such an interface and it may be helpful to provide meta-information stating the connector capabilities in order to dynamically guideline the code generation for a certain connector.

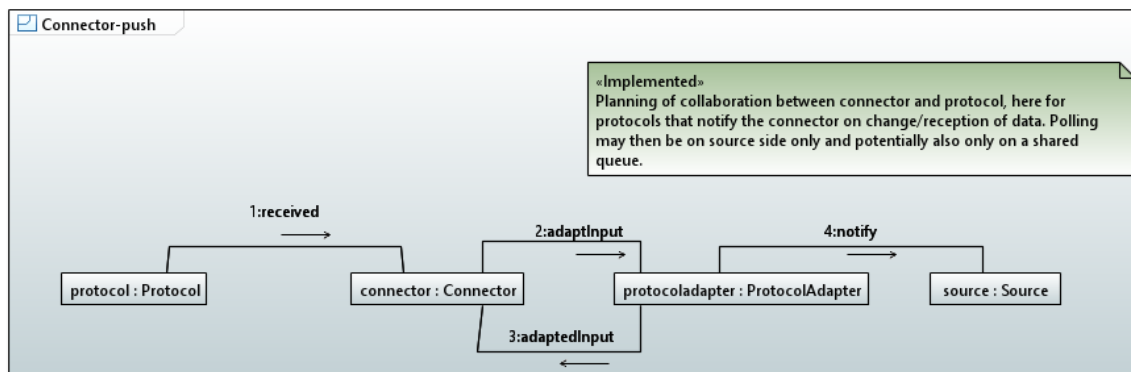


Figure 31: Event-based connector and push-based protocol-adaptation.

Moreover, connectors may differ in their data provisioning style. For performance reasons it is desirable to utilize **event-based ingestion**, i.e., the underlying protocol or information model informs the connector about new or changed data. Message passing approaches like MQTT or information-model based approaches like OPC UA provide such events. In this case, as illustrated in Figure 31, the “Protocol” notifies the Connector about new data. In turn, the Connector consults the ProtocolAdapter to translate the external data into an application-specific type, which, dependent on the “Protocol” capabilities, can be done in terms of payload translation or by querying the abstracted model of the “Protocol” (not shown in Figure 31). When the data is translated, the respective instance is passed on to a registered streaming Source in asynchronous manner. For the outbound direction (not shown in Figure 31), the Source ultimately receives the data as a stream and calls the Connector upon a received data item, which then consults the ProtocolAdapter in the backward direction ultimately leading to a send/write command on Protocol.

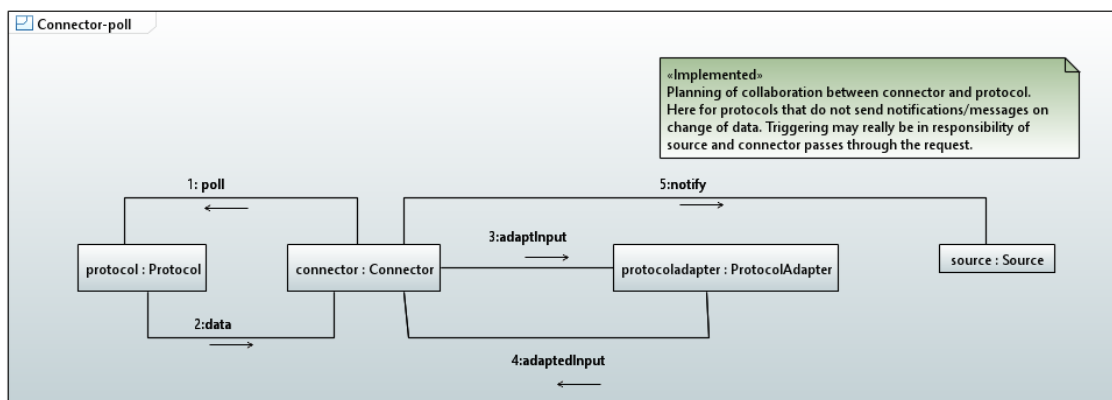


Figure 32: Poll-based connector and subsequent protocol adaptation.

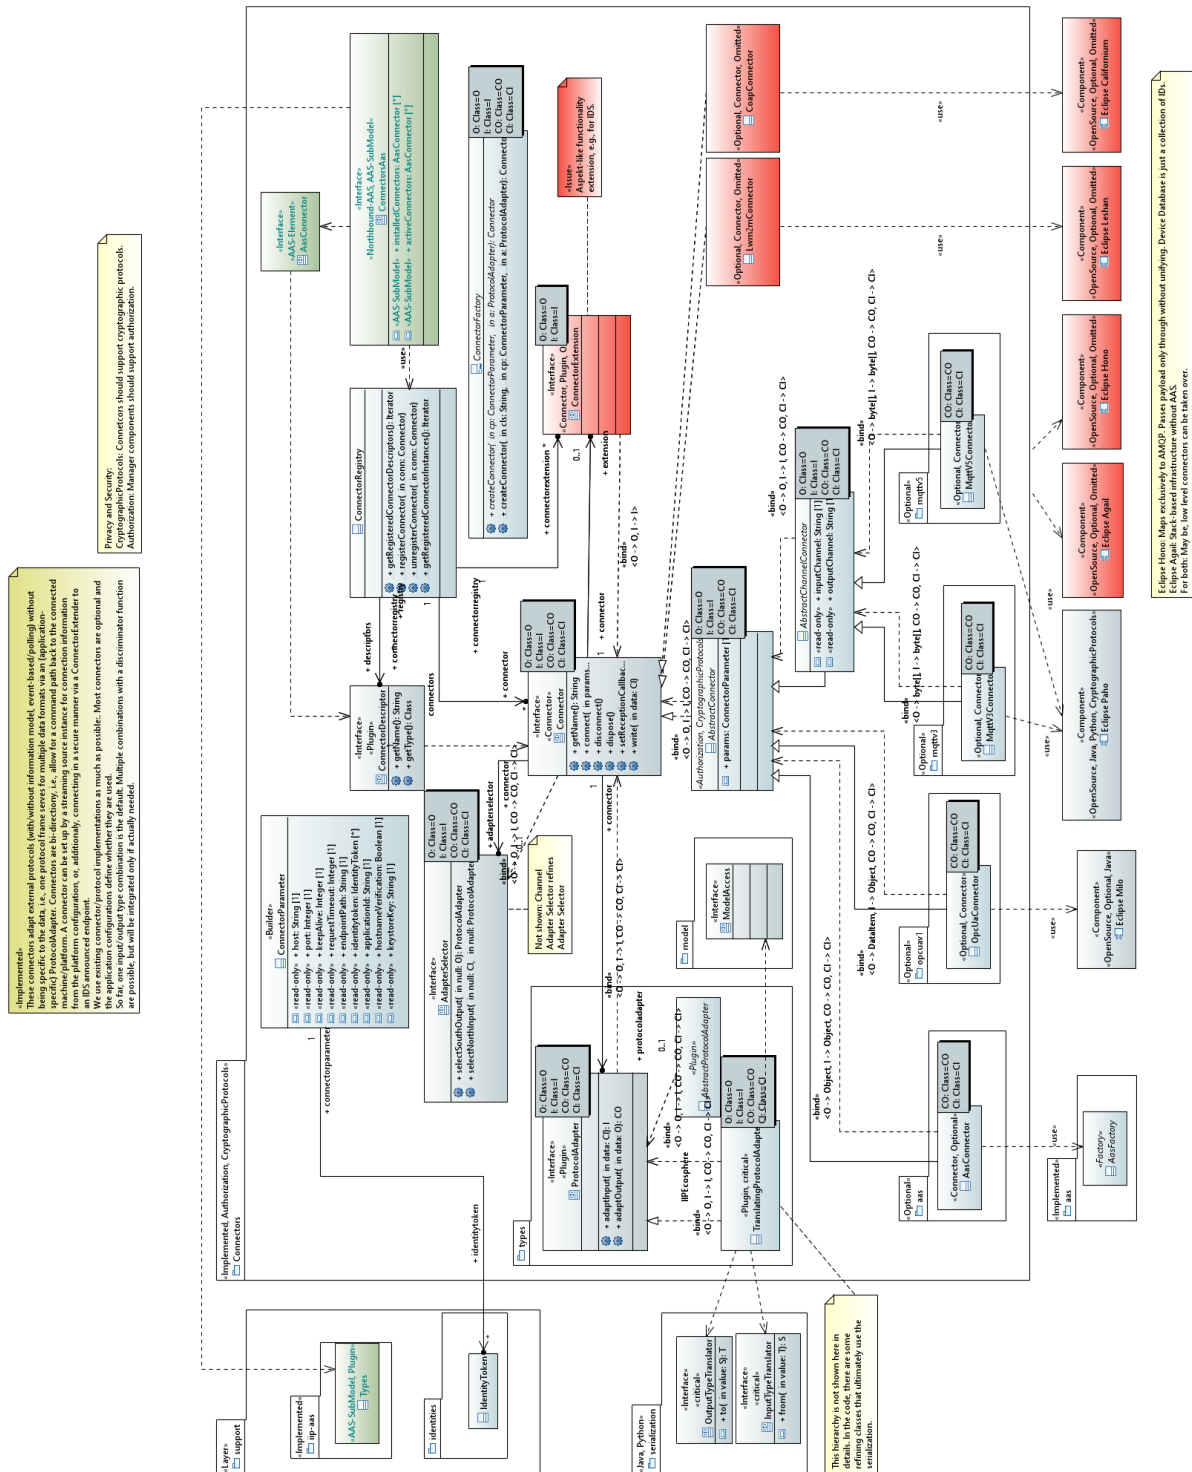


Figure 33: Overview of the Connectors Component (comments partially cropped).

Protocol implementations not offering such notifications are subject to **polling**. One example here is the current BaSyx implementation of AAS. In the version that we currently use, no events are provided (BaSyx plans for events earliest end of 2020⁷⁷). In its basic form, polling repeatedly obtains information from a source and ingests the information regardless whether the information changed in the meantime. This may be intended, e.g., to realize equidistant input. However, if not desired, it can also unnecessarily allocate transport resources. To avoid this, since version 0.5.0 of the platform, a caching

⁷⁷ If the required notifications are available, the AAS machine/platform connector can be extended to support event-based ingestion.

mode can be defined for a connector instance. The mode indicates whether there shall be **no** caching, i.e., ingestion of all received data, ingestion only if **hash** codes are different or ingestion if the contents of the data is not **equal**. While a comparison based on hash codes is typically faster than a comparison for equality, it may also fail if hash codes accidentally are the same for equal data items (hashing collision).

As illustrated in Figure 32, the IIP-Ecosphere Connector then actively (based on connector settings) polls information from the “Protocol”. As before, the Connector consults the ProtocolAdapter and notifies the registered Source about the data to be ingested. The outbound direction works as discussed for event-based ingestion. Realizing the polling cycle in the Connector rather than the Source allows for connector-specific polling strategies as well as for a uniform interface towards the stream-based data processing in the IIP-Ecosphere platform.

While event-based injection and polling may appear to be an alternative choice, a Connector may, if feasible, implement both alternatives and let the user (via the setup/platform configuration) decide about the desired approach. In particular, connectors for protocols based on information models may support both forms (such as OPC UA). Figure 33 presents an overview of the main classes in the Connectors Component of the IIP-Ecosphere platform. The component consists of:

- The Connector interface in the middle of the figure representing a platform/machine connector. Connectors based on an information model shall exhibit a ModelAccess instance to interact with the information model. The Connector interface defines four template parameters, consisting of the data types accessible from the platform, i.e., CI for input into the connector and CO for output produced by the connector, and the data types for the handshake with the underlying protocol implementation, i.e., I for input into the protocol and O for output issued by the protocol. A Connector can be connected as specified in the ConnectorParameters and security settings like IdentityToken or certificates. When connected, received data of type O is passed through a ProtocolAdapter and an interested party is informed through a ReceptionCallback (from the Transport Component) in terms of a data object of type CO. Via the write method, data of type CI can be passed in, is translated by the ProtocolAdapter and handed as an instance of I to the underlying protocol. Finally, a Connector can be disconnected or, ultimately, disposed. So far, we plan for a single distinct pair of input/output types. If heterogeneous types shall be covered, we see two alternatives: 1) Mapping the alternative types as alternatives into an umbrella type. 2) Using a discriminator in terms of the AdapterSelector.
- The TranslatingProtocolAdapter is a default implementation of the ProtocolAdapter and relies on type translators, i.e. InputTypeTranslator and OutputTypeTranslator defined by the Transport Component. The ProtocolAdapter and its related classes will be detailed below. In particular protocol adapters to information models have a relation to a ModelAccess instance, which allows the type translation to interact with the model.
- The AbstractConnector provides a basic implementation, e.g., for handling the ReceptionCallback, for utilizing the ProtocolAdapter, etc. leaving just methods open that are protocol specific. The AbstractChannelConnector specializes the AbstractConnector for channel-based protocols such as MQTT and, in turn, requires a specialized protocol adapter (as we will detail below).
- The ConnectorExtension may add additional capabilities to a connector, e.g., IDS support. The IDS reference architecture model introduces the concept of Trusted Connector. Such a connector extends the security features of Base Connector. An IDS connector generally focuses on security and delivers a trusted platform, incorporating several mechanisms such as identity and trust management for authentication, trustworthy communication based on encrypted

connections. Instances of trusted connector allow the remote integrity verification to ensure the integrity of the deployed software before granting corresponding access to data. Such connectors guarantee a controlled execution environment for data services [22].

- The ConnectorRegistry collects information about installed and used connectors. Installed connectors are registered through an instance of ConnectorDescriptor upon infrastructure startup (in Java through JSL) with the ConnectorRegistry.

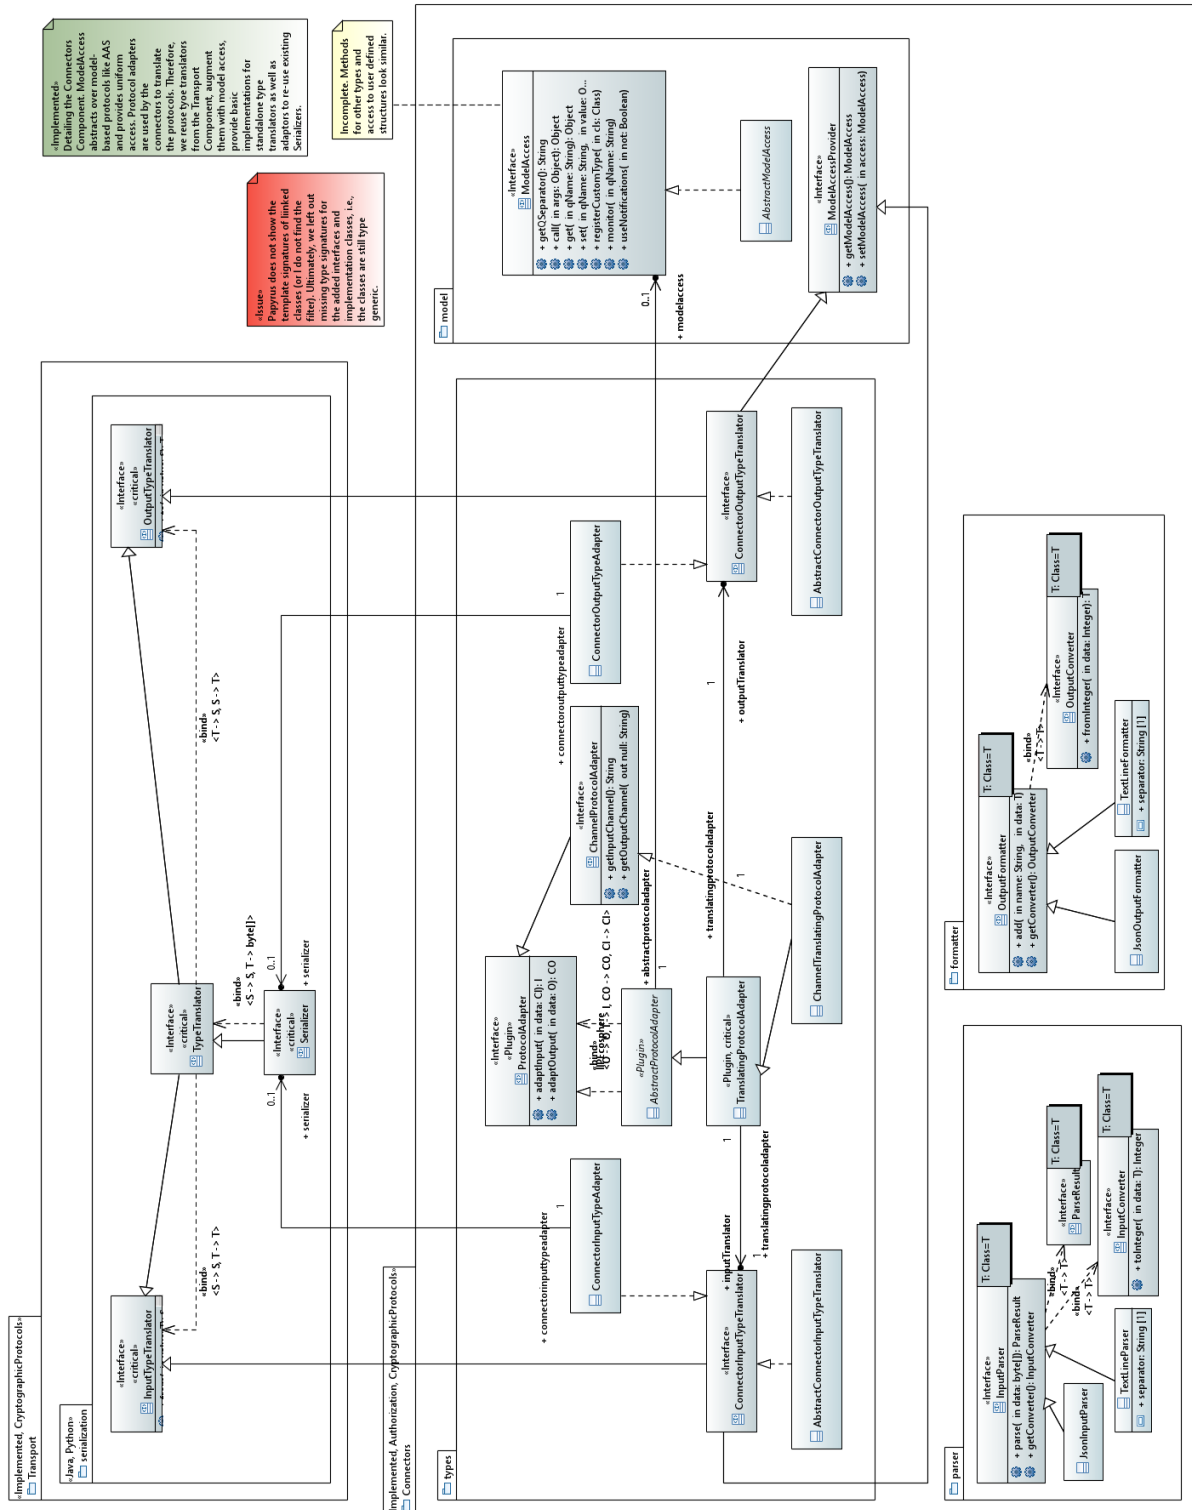


Figure 34: Model Access and Protocol Adapter in the Connectors Component.

- The information provided by the `ConnectorRegistry` is also the basic information to be presented in the AAS of the Connectors Component. Further, selected capabilities of the connectors are made available through the `installedConnectors` sub-model of the platform AAS. Created connector instances register themselves upon connect/disconnect with the `ConnectorRegistry`, which in turn leads to an update of the `activeConnectors` sub-model, i.e., connected connectors appear as sub-model elements and disconnected connectors are flagged as inactive⁷⁸. Further, connector instances provide access to their input/output data types by referencing to the respected sub-model elements in the `Types` sub-model (see Section 3.5). Ultimately, connector instances link to their descriptors in the `installedConnectors` sub-model to indicate their origin and capabilities.
- The `ConnectorFactory` is a proxy to dynamically create the most appropriate connector instance if there are alternatives, e.g., MQTT v3 and MQTT v5. Such a `ConnectorFactory` takes the `ConnectorParameters` and may decide on the supplied device service information (if present) on the actual connector type. An implementing connector artifact may rely on existing connector implementations (as we do for MQTT). The configuration model (see Section 6) shall define a specific type representing the `ConnectorFactory` and its artifact to ease its application.

Currently, four specific (optional) connectors are realized in terms of individual components extending the Connectors Component. These are the generic `AasConnector` for integrating external AAS into the platform (based on the `AasFactory` from the Support Layer⁷⁹), the `OpcUaConnector` for OPC UA 1.04 (based on Eclipse Milo) as well as two payload-based MQTT connectors, one for MQTT v3 and one for MQTT v5, also based on Eclipse Paho akin to the Transport Component. Each of these protocols bind the known template parameters of `Connector` as needed, all leaving CI and CO unbound as these are application-specific types to be defined when instantiating the respective connector (and providing a matching `ProtocolAdapter`). Optional TLS encryption support (R40) is realized for both MQTT connectors, the OPC UA connector and the AAS connector (if the BaSyx implementation is used, the AAS registry remains unencrypted by default, the AAS server can optionally be executed with TLS). These approaches/protocols have been selected due to the required mandatory support for BaSyx (R7) as well as R14a explicitly mentioning OPC UA and MQTT. All connectors are based on TCP/IP networking (R14b). However, R14b does not clarify the protocol and the wire format, but some TPC/IP based protocol can be realized using the structures defined in the Transport and the Connectors Component. Similarly, Bluetooth LE was mentioned in R14c without further details. As the Connectors Component is designed to be open, further connectors can be added easily.

We will now detail the `ModelAccess` and the `ProtocolAdapter` approach illustrated in Figure 34. Some approaches like OPC UA or Asset Administration Shells (AAS) are based on an information model, OPC UA even with user-defined custom types. Accessing this model in a uniform manner is a key requirement for simplifying the generation of application-specific code for the connectors. This is specified by the `ModelAccess` interface, which allows to read/write properties (based on a hierarchical naming scheme to be interpreted in the context of the underlying protocol), to call operations, and to register (the implementation counterpart of) custom types. The specific `ModelAccess` instance of a `Connector` can perform instance translations between value instances of the model and the actually used internal type in the platform/application datatypes. `ModelAccess` provides also opportunities to establish monitors on the underlying information model of the protocol,

⁷⁸ So far it seems that no elements can be removed at runtime from an AAS, potentially to not render references among them illegal.

⁷⁹ While BaSyx is the default implementation for IIP-Ecosphere, this connector provides the possibility to define the individual instance to be used, i.e., individual instances for specific connections may use other factory instances than the default one.

i.e., to be notified on specific changes, as well as to register programming language counterparts of custom types defined in the model.

While all methods can be implemented for OPC UA, not all methods are currently meaningful for AAS or at least the version of BaSyx that we are using, i.e., some capabilities may not be supported which can be indicated in the meta-data of a Connector. The use of the abstracted model access is supported by `AbstractModelAccess` providing a common basic implementation. For payload-based protocols such as MQTT, implementing the `ModelAccess` interface is not required.

As illustrated on the left side of Figure 34, several further interfaces and classes are defined to support the type translation. These classes shall provide flexible support for applying the type translation mechanism and even to utilize existing serializers if applicable. If required, a Connector may be also implemented from scratch, i.e., without a `ProtocolAdapter` or using a `ProtocolAdapter` that is not based on `TypeTranslator` from the Transport Component. The `AbstractProtocolAdapter` is a default implementation providing access to the `ModelAccess` instance of the connector. The `TranslatingProtocolAdapter` is a refined default implementation of `ProtocolAdapter` using two type translators, one for the translation of information provided by the protocol, i.e., types O to CO, and one for the translation in counter direction, i.e., types CI to I.

As discussed above, the role of the `ProtocolAdapter` and the involved type translators changes when the underlying protocol is based on an information model rather than on payload transport. For payload transport, the target communication channels are needed, which are introduced by the `ChannelProtocolAdapter` and its default implementation `ChannelTranslatingProtocolAdapter`, an extension of the `TranslatingProtocolAdapter`. In contrast, for an information model-based protocol, also the `ModelAccess` instance must be made available to the type translators as well as further initialization work such as defining the polling mode must be performed. This is introduced by the two refining type translator interfaces, namely `ConnectorInputTypeTranslator` and `ConnectorOutputTypeTranslator`, both with a corresponding basic implementation.

A further aspect is covered by the types package on the left side of Figure 34. The direction of type translation of `Serializer` in the Transport Component is opposite to the direction for a `ProtocolAdapter` in the Connectors Component, i.e., `Serializer` instances do not fit directly into these. In some cases, it would be convenient to use already defined `Serializer` instances. To facilitate this reuse, we introduce the `ConnectorInputTypeAdapter` and the `ConnectorOutputTypeAdapter`, which both take a `Serializer` as input and make it usable in the context of a Connector.

At the bottom of Figure 34, the two packages `parser` and `formatter` depict how the platform can handle binary input or output of channel connectors in a generic and open manner. The platform offers certain default input parsers or output formatters, respectively, currently JSON and data formatted by separators, e.g., tabulators or commas as we found in some application cases. Parsed data is supposed to be mapped to record types defined by the user in the configuration model. After parsing, data is either accessible by hierarchical names, i.e., in terms of nested data types, or in positional manner along a deep traversal of the defining data type. In the opposite direction, data is handed in deep traversal sequence to the data formatter, which produces the respective format in terms of binary data to be passed to a channel connector. Individual data values can be read in typed fashion from an input parser via an associated input converter. In the opposite direction, the output converter takes typed data and converts it into the output format. Both mechanisms are utilized by the generated connector integration that we discuss in more detail in Section 6.

3.6.3.3 Validation

The functional validation of the Connectors Component and the specific connectors realized as extensions happens through regression tests. Therefore, we follow the same basic idea as explained for the Transport Component in Section 3.6.2.3, i.e., we set up a corresponding protocol server/broker and cause an information shortcut between server side and test code. The test code produces protocol output data (of type O) either by modifying the underlying information model (event-based ingestion, polling) or by sending respective payload. The connector under test translates the data and issues an instance of type CO to a `ReceptionCallback` in the test code, which turns the information into an instance of CI and writes it back into the connector. The respective information must occur on the protocol side and can be analyzed and asserted by the test code. Also, these regression tests are subject to Continuous Integration.

Further functional tests have been performed in the context of the IIP-Ecosphere platform use case studies, e.g., in the context of partners such as Lenze, Phoenix Contact or the VDW. In these validations, the default IIP-Ecosphere connectors were complemented with handcrafted serialization or type translation plugins to test the expected intake of use case specific formats (e.g., OPC UA or JSON via MQTT) and, where applicable, connected to the infrastructure of the partners. Moreover, the connectors are subject to automatic, model-based integration via the configuration model and the platform instantiation (cf. Section 6). For validation, the situations of the partners from the platform use case studies have also been modeled in terms of the IIP-Ecosphere platform configuration, the connector service integrations have been generated and the intake has been validated. We conducted initial performance analysis on the serialization mechanisms and the generated connectors. The comparison of serialization mechanisms indicated performance peaks up to factor 10 for different JSON libraries in the context of the use case studies. The generated connectors are (nearly) as fast as the handcrafted ones, sometimes sacrificing a bit time for a generic, open integration as well as a more schematic model-based integration.

3.6.4 Requirements Discussion

Finally, we review in this section the realization of the most relevant requirements for the Transport and Connection Layer. The results are summarized in Table 9.

Table 9: Review of realized⁸⁰ requirements for the Transport Layer (based on Table 2, Table 3 and Table 5).

Requirement	Summary
R1	Support for protocol/streaming extensions of different vendors based on different technologies.
R2	Used standards: MQTT, AMQP, OPC-UA, AAS
R3	Virtual platform: Container integration possible, communication with underlying platforms possible but not responsibility of this layer
R4	Design based on components and services
R5	Eclipse Paho, Rabbit MQ, Eclipse Milo, Eclipse BaSyx, Spring Cloud Stream, more for testing
R6	Open for optional/commercial components (transport connectors, serializers, machine/platform connectors, protocol adapters, etc.)
R7	Basic information on transport as well as available/active machine/platform connectors is provided. More information regarding supported protocols or broker may follow.

⁸⁰ In the requirements review tables, „realized“ refers to implemented in terms of functionality, committed into the IIP-Ecosphere GitHub repository, tested and integrated with the platform functionality. Text in *italics* refers to missing functionality, i.e., entries that are partially formatted in *italics* typically indicate partial realization. Work in progress or incomplete/non-integrated realization may be excluded from platform releases.

Requirement	Summary
R8	On component/extension level: Alternative TransportConnector and Spring Cloud Stream binder components without cross-dependencies among the protocols (except for re-use in testing). Optional platform/machine connectors without cross-dependencies among the protocols (except for re-use in testing).
R9	<i>Level 1:</i> By auto-reconnect mechanisms of the protocol implementations and of the streaming library <i>Level 2:</i> By monitoring the service execution and restarting services if needed (see ECS runtime in Section 0) <i>Level 3:</i> By explicit fallback, i.e., hot-standby replication of services, multiple connected broker installations and dynamic stream rerouting. Level 3 is supported by the selected streaming library, <i>but not realized in this release.</i>
R10	Soft realtime processing (<100 ms) for production-critical functions feasible (see Sections 3.6.2.3 and 3.6.3.3, excluding services on this level)
R11	Documentation (also in terms of this section), extensive code documentation with JavaDoc, generation subject to Continuous Integration and Maven deployment
R12	<i>Not implemented in this release.</i>
R13	Connectivity to other actors via standardized and open protocol integration
R13a	I4.0 devices: Via standardized/open protocol integration and flexible wire formats.
R13b	I4.0 platforms: Via standardized/open protocol integration and flexible wire formats.
R13c	Other IIP-Ecosphere platform instances: Supported via implemented AAS connector, optionally using IDS functionality. Supported by standardized transport protocols and flexible wire formats.
R14	Open and flexible connector integration and default connectors for platform internal messaging and stream transport.
R14a	MQTT is supported. OPC UA PubSub is currently not feasible on this layer due to missing implementations. Services and Connectors Layer (cf. Section 3.7) will take OPC UA into account.
R14b	TCP/IP support by all implemented connectors. Further, plain TCP/IP protocols with flexible wire formats possible through extensibility of discussed components.
R14c	<i>Future: Bluetooth LE may be supported through additional connectors if required.</i>
R15	Uniform connectors through two main interfaces as well as type transformation / serialization interfaces.
R16	Open and flexible connector integration shown for five protocols. Extension by further protocols from partners and externals possible, supported through Open Source development GitHub.
R17	Distribution of connectors to devices by considering the Transport Layer as re-usable, deployable component only with dependencies to selected protocol implementations.
R17b	Management of connectors by platform through design management classes, exhibited by respective AAS (R7)
R17c	Parameterizable connectors through parameter objects and connector plugins such as the ProtocolAdapter.
R18	<i>Connectors with specific security mechanisms are not part of this release.</i>
R19	Minimal number of internal wire formats through common type-safe data serialization while enabling application-specific data types. Feasible wire formats can be selected through serialization implementation/generation.
R19a	Example input formats (southbound), covered by application-specific types and generic serialization. 50 values per data item feasible see Sections 3.6.2.3 and 3.6.3.3.

Requirement	Summary
R19b	Example input formats (northbound), covered by application-specific types and generic serialization.
R19c	Restful APIs with JSON/XML through AAS implementation and wire format as argued in this section.
R19d	Example output formats (northbound), covered by generic serialization
R19e	Output data clocked in 5 s intervals possible (see Sections 3.6.2.3 and 3.6.3.3, excluding services on this level)
R19f	Via the <code>TypeTranslator</code> , configuration and code generation
R19g	Not part of this layer, supported at least through <code>TypeTranslator</code> , configuration and code generation
R20	Supported by streaming library, to be realized by glue code generation (through the configuration model).
R21	Low impact on data throughput (see Sections 3.6.2.3 and 3.6.3.3, for existing connectors, excluding services on this level)
R22	Platform data throughput of 500 GBytes per year (see Sections 3.6.2.3 and 3.6.3.3)
R28	Machine pulse of 8 ms feasible (see Sections 3.6.2.3 and 3.6.3.3)
R35	OT sampling frequency of 2 ms does not apply to the IT side.
R38, R39, R41-R68	<i>Several security and privacy mechanisms that are introduced in Figure 20 are used to ensure introduced security and privacy mechanism. These mechanisms can be added to different layers, phases, and abstraction levels of design. Security mechanisms are indicated in the architecture but not part of the implementation of this platform release.</i>
R40	The default transport protocol plugins/binders as well as the MQTT machine connectors provide optional support for TLS, <i>but so far no RBAC mechanism is integrated.</i>
R91	7 GByte per hour not validated on this level (see Sections 3.6.2.3 and 3.6.3.3, excluding services on this level)

We conclude, that most of the basic requirements for this layer are already implemented. Advanced functionality as well as security and data protection mechanisms (although prepared through respective abstractions) are subject to one of the next releases.

3.7 Services Layer

The Services Layer introduces the basis for deployable services, i.e., their interfaces, data flows, monitoring support, management and AAS representation. We separate this layer into two major components, one component to control/manage service instances and a second providing a unified execution environment for services. We start with a discussion of the terminology and background in Section 3.7.1 and detail then the requirements for this layer (Section 3.7.2). In two further subsections, we turn then to the two major components of this layer.

The service management component is generic and can be realized in the same way for all services. Due to the overall vision of IIP-Ecosphere to support easy-to use AI methods in intelligent production, AI functionalities shall be realized in terms of services (“AI services”, named as AI toolkit in [13]). Nowadays, AI is typically realized using various programming languages, in particular Python (R113 names Python and Java). We support this in terms of a language-specific execution environment supporting a unified integration and easing the development of services for the IIP-Ecosphere platform. In Section 3.7.3, we discuss the Service Execution Environment for Java and Python.

The Control and Management component (Section 3.7.4) is closely related to ECS runtime and acts as control interface for the platform to take command over services running on certain devices. Control operations are, e.g., starting, stopping, reconfiguring or updating services. These operations are offered through an AAS, which also provides access to runtime monitoring information for individual services. Specific operations involve multiple services, such as switching among equivalent services or migrating services among resources, where the control and management component is responsible for the orchestration of such operations.

While we briefly discuss the validation of the individual components at the end of the respective section, we review the requirements for this layer in Section 3.7.5.

3.7.1 Terminology and Background

In this section, we briefly introduce our notion of the term service and discuss the bigger picture, where service implementations are supposed to originate from.

Several notions for services are used, ranging from web services to microservices. In the IIP-Ecosphere platform, a **service** is (a thread in) a process implemented in any programming language. A service receives data and produces data in continuous stream-based manner. Therefore, the platform keeps a service alive and runs it continuously. Input and out data types are defined and their correct composition shall be controlled through the configuration model and the subsequent code generation (cf. Section 6). Data handling can happen synchronously, i.e., an input item is turned directly into zero or multiple output items, or asynchronously, i.e., the service receives data and produces data at any time later if at all. A service indicates its state (R4c), meta-information (R4b), name, identification, version, kind/category as well as the typed input- and output connectors (R4a). Moreover, it allows for certain runtime operations such as passivation, migration, runtime switch to an equivalent service, reconfiguration or (re-)activation. Services are typically connected to other services of different kind, ranging from source over transformation to sink services. A specific service kind is the probe service, with inspects data for monitoring, but passes the data through (typically) without modification. Instantiated connections between services are called data paths/relations [41] and shall be defined as part of the applications in the platform configuration.

Services can be **classified** into **platform-provided services** and **application-specific services**. A platform-provided services (for short platform service) is a service that is shipped with the platform. Typically, a platform service is generic and parameterizable so it can be applied in various settings. As services usually were not realized for the IIP-Ecosphere platform, the integration of such a service and the form how to customize the service, e.g., with respect to required input/output data formats, the data transport to/from the service, the passing of parameters, etc. is typically rather service specific. In contrast, an application-specific service is (initially) designed for the use within a specific application, i.e., input/output datatypes as well as parameterization are fixed and determined by the application context. Moreover, **hybrid services** may occur, typically generic platform services that can operate without application specific code, but which can be customized through add-ons or plugins, which then turn an instance of the hybrid service into an application-specific service. Thus, hybrid services can ease the realization of application-specific services and still act as platform-provided services.

Further, it is important to answer the question “**Where do services come from?**”. Details of the mechanisms will be introduced later, in particular in Sections 3.11 and 6. However, a coarse picture may already be helpful here. Service types are specified in the platform configuration, in particular through their technical information, their meta-information and the input/output datatypes. Also, the relations among the services in terms of application-specific service meshes are defined in the platform configuration. The platform instantiation/code generation either integrates platform-provided services into the realization of such a mesh or, in case of application-specific services, turns this

information into service interfaces. Moreover, the code generation creates support artifacts such as data classes, data serializers or basic service implementations (for any relevant programming language, e.g., Java or Python). Further, the instantiation process binds the service (interfaces) with service/glue code to the selected streaming engine. For application-specific services, the binding happens through dynamic class loading. First, this allows for code generation even if there is no actual service implementation, e.g., when the model is built up and in a first stage only service interfaces shall be generated. Further, it eases packaging with reduced dependencies as class loading cannot fail if dependencies for services that shall not be executed are intentionally missing. Last, dynamic class loading also eases dynamic exchange of compatible service implementations, e.g., when adapting the service execution. Dependent on the service configuration, data may be handled synchronously or asynchronously. As part of the generation process, also the service descriptors required by the Service Control and Management component or the Spring application setup including the service wiring are created. Here also information for AAS nameplates is taken from the configuration and turned into respective implementation artifacts to build up the related AAS and submodels, e.g., one vendor AAS per service type known to the platform.

3.7.2 Requirements

In Table 10, we summarize the specific requirements for the Services Layer discussed in [13]. The notion of a service is cross-cutting, i.e., it occurs in many topic areas in [13] and, thus, a summary of all relevant requirements is important for the design and realization. Besides these functional requirements, we must also consider the decisions made so far, i.e., that services may offer a two-folded communication: 1) communication at lower pace for commands, status and quality properties via AAS and 2) soft-realtime communication via streams whereby the stream-integration shall be generated and flexible in order to allow for an exchange of the streaming approach. This is in particular important for monitoring (R4b, R4c, R4e, R4f, R133) of runtime properties and the runtime stream management, in particular to start, stop, connect (R20), update (R135), configure (R32), adapt (R69 and R31c, see also dynamic service selection in [13]) or dispose (R134c) services on demand. To be integrated in a flexible manner, monitoring and service management must be realized based on explicit interfaces, so that an exchange of the implementations becomes possible. If feasible, existing interfaces shall be utilized.

In the default stream processing approach in IIP-Ecosphere, i.e., Spring Cloud Stream (see Section 3.6.2.1), the micrometer⁸¹ interface is used to exhibit monitored information in HTTP/REST style. Moreover, Spring provides specific management capabilities for Spring Cloud based service applications, e.g., to start services in individual processes. As micrometer is supported by several (commercial) monitoring tools, it appears to be a valid choice as monitoring interface, which, however, must be integrated with AAS (R7). For the stream management, it makes sense to reuse existing functionality from the Spring Cloud ecosystem, e.g., the Spring Cloud Deployer. It is important to make this functionality optional and to enable it when Spring Cloud Stream is selected as stream processing engine of a platform instance. However, most of the Spring management providers target rather specific (non-edge cloud) environments, so we will rely on the so called Local Deployer and integrate it as optional extension of the Service Control and Management component.

Moreover, the Service Layer must set the scenes for the management of heterogeneous service implementations (R113), including platform services that are more likely to be realized in Java or as Java interfaces to service implementations of underlying frameworks or platforms.

⁸¹ <https://micrometer.io/>

Table 10: Requirements for the Services Layer (excluding configuration, storage services, not repeating Table 2 or Table 3)

Requirement	Summary
R4a	Components/services must be described with their interface (input, output)
R4b, R131b	Components/services must be equipped with meta information (version, categorization)
R4c	Components/Services must have a queryable state
R4d	The execution of the services must be supervised
R4e	Service monitoring shall be parameterizable
R4f	Service monitoring shall be realized by application-specific services
R20	Application-specific data paths
R20b	Data paths can have properties/parameters
R20c	Data paths shall be managed by the platform
R29c, R70, R122f	Services shall describe their own quality properties and functions as AAS
R31	Container shall contain only the required components/services
R31a	Required components must be specified in the configuration model
R31b	Containers can contain optional components
R31c, R69	Alternative services for one task, even dynamic exchange of (alternative) services at runtime
R32	Configuration of services via parameters
R39	Personal data processing only for authorized users
R41	The security mechanisms shall be integrated with common directory services
R42	Further safety mechanisms must be configured uniformly
R46	Collection of personal data must be for specified, clear and legitimate
R47	Avoiding the processing of personal data as much as possible
R48	The platforms should not store data for longer than necessary
R49	Process personal data adequate and relevant to the legitimate purposes
R52	Store personal data in a structured, common and machine-readable format
R67	Capture and classify generated cookies or similar identifiers stored on end devices
R73a-f, R79	Supported datatypes: structured, heterogeneous time series, unstructured data, labeled data, meta data/data schema
R113	Support for different programming languages, e.g., Python
R132	Platform-supplied and application-specific services shall be supported
R133	Runtime support for applications (and the services an application consists of)
R133a	Status of services
R133c	Support for changing the status of services
R133d	Detection of failure states and functions to mitigate failures
R134c	Removal/disposal of services
R135	Update of applications (and the services an application consists of)

3.7.3 Service Environments

In the IIP-Ecosphere platform, the service environments provide implementation and execution support for services realized in different programming languages. Java services and non-Java services are integrated differently into (a Java-based stream-based) service execution engine. While Java services can be directly called, non-Java services are executed as individual processes and receive their control commands and data via inter-process communication/network, in particular an AAS command server. It is desirable to use a single server here, e.g., the “AAS command server”.

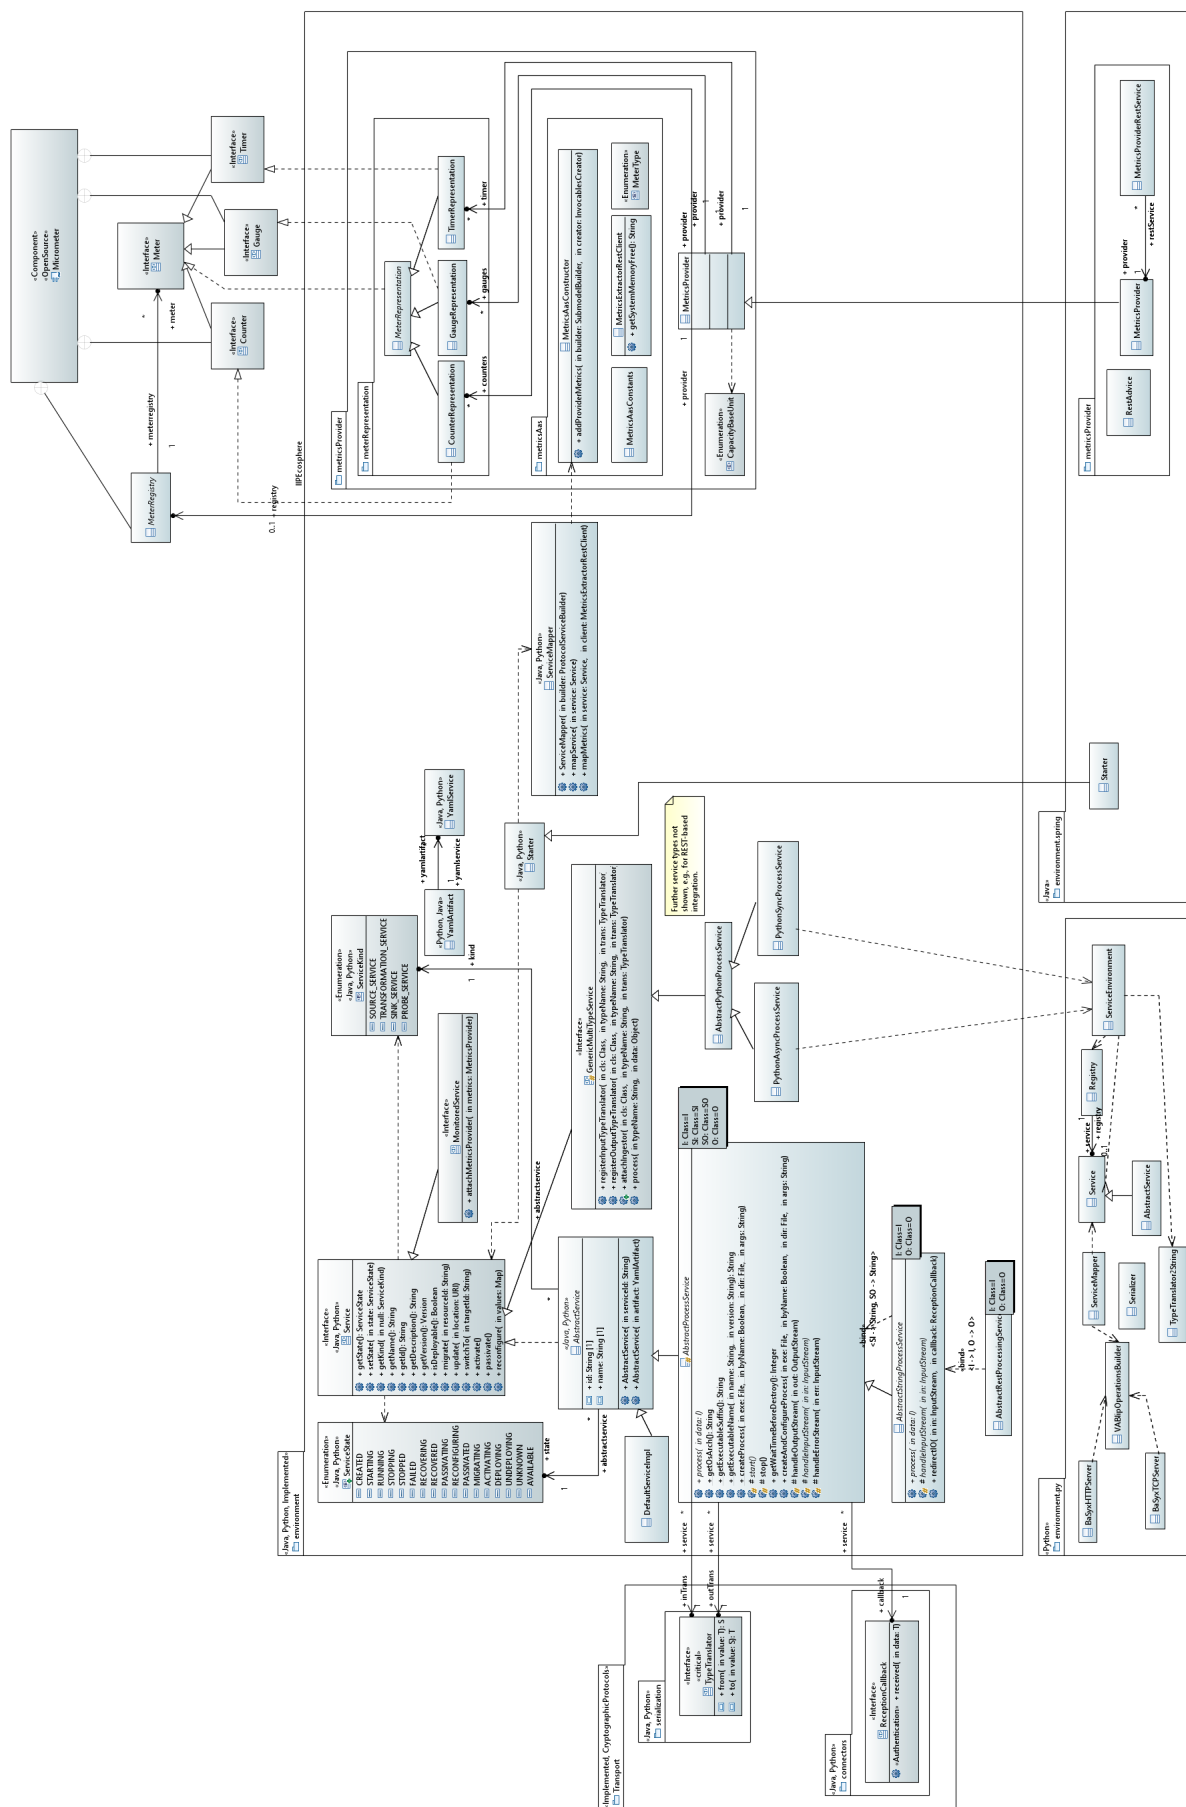


Figure 35: Design of the Service environments.

As the configuration modeling approach does not target the modeling of behavioral code rather than configuration options and their interdependencies, the actual implementation of services cannot be produced in this step. Such an approach would require a high effort in modelling, in particular (a combination with) modeling approaches that are better suited to represent algorithm behavior, call sequences etc. We do not target such complex models as the realization of functional services shall ultimately be a human activity, e.g., by a data analyst, AI expert, etc. and shall be integrated with their work. Thus, the interfaces produced by the code generation are input to a manual implementation process and interrelate in particular with later phases of the data analysis lifecycle. For this purpose, the code generation produces interfaces, data classes and serializers for all supported languages, i.e., currently Java and Python. The resulting code shall lead to versioned code artifacts (Maven artifacts, Python code with “python” as classifier) and be deployed into a repository. For the aforementioned deployment descriptors and the automated creation of containers (not part of this release) for the service execution, also non-Java dependencies, e.g., to AI frameworks are relevant. The respective artifact information is specified for the individual services in the configuration model and the packaging of the application artifacts at the end of the code generation process integrates the service code with service descriptor and application setup. Ultimately, the application artifacts are made available through an own repository mechanism (called the Service store in [41]) so that either the automated generation of containers or the Service Management and Control component on the target resources can obtain the service code and install/execute it.

3.7.3.1 *The Java Service Environment*

Figure 35 illustrates the concepts and relations of the service environments, in particular the **Java Service environment**. The central package (environment) represents both, the basic service environment for Java and Python. This package (on the left side of Figure 35) defines the Service interface with all operations discussed for an IIP-Ecosphere Service in Section 3.7.3. The states from Figure 36 are represented in terms of the ServiceState enumeration, the four main kinds of services in terms of the ServiceKind enumeration. The service environment also provides support for service parameters, i.e., the customization of generic services through values that are determined upon service start or may be changed at service runtime.

A service realization is free to fill the service meta-information as desired, e.g., through code generation or by reading the information from a file. As the Service Management and Control component relies on service deployment descriptors, one obvious approach is read out the relevant information from that descriptor. As these descriptors are given in YAML format, the two classes YamlArtifact and YamlService are part of the service environment to read and represent that information. It is important to recall that we need here only a part of the information in the deployment descriptor, e.g., the technical information on how to transfer network ports or how to start a Python process are not required. Thus, the two classes represent only the relevant information and the YAML parser is commanded to ignore all further information. In turn, both classes can be used as a basis to realize the parsing of the deployment descriptor of the service management and control operations in Section 3.7.3. For this purpose, parts of the (Java) service environment are imported into the Service Management and Control component and used there.

Further abstract classes provide basic mechanisms for realizing services, e.g., through operating system processes or using the REST protocol (on device-local networks). We just mention some examples here. AbstractProcessService provides the abilities to create an operating system process through a standardized naming scheme, to manage the process instance via the service status, including activation and passivation, and to customize the console input/output streams. The AbstractProcessService defines four template parameters, namely the received input data type (I) from the perspective of the platform/application, the input data type (SI) of the implementing process, the output data type (SO) of the implementing process and the output data type of the service

(O) from the perspective of the platform/application. An `AbstractProcessService` requires two `TypeTranslator` instances, so that incoming data can be translated into a format that can be handled by the implementing process and, further, that the output of the implementing process can be re-ingested into the platform data streams. Here, we rely on the `TypeTranslator` interface from the Transport component (cf. Section 3.6.2) rather than the `Serializer` interface as the format to be passed to the implementing service could be JSON, XML etc. and must not necessarily be binary data. Moreover, the `AbstractProcessService` requires a `ReceptionCallback` (the interface from the Connectors component in Section 3.6.3), so that data can be processed asynchronously. Currently, the classes do not (yet) support synchronous processing. As a refinement, the `AbstractStringProcessService` forms the basis for process-based services that communicate through a string format via console streams with the implementing process, e.g., JSON. This fixes two template parameters, namely `SI` and `SO` to `String`, also allowing to provide a basic realization of some inherited abstract methods, e.g., how to receive data from the implementing process. The `GenericMultiTypeService` supports multiple input/output types in a generic manner. The Python integration services (for the Java side) utilize these capabilities to integrate Python services via the Python service environment (see below). Moreover, basic classes for REST-based service integration do exist, e.g., the `AbstractRestProcessService`.

One generic service that is provided by the Java Service environment is the `TraceToAasService` (not shown in Figure 35). This (sink) service provides/contributes to an AAS used as application endpoint. However, it does not show the received data rather than optional trace messages sent by the individual services in an application. Trace messages can be enabled for debugging or for demonstration. These messages carry their origin, the action as well as an action-specific payload, e.g., the received data. The service collects all trace messages and displays them for a given time frame in its trace submodel. Moreover, the service is intended to act as a hybrid service, i.e., it can be used as basis to implement an application-specific service, which then may display processed data in an application manner or which may provide operations to be called by a device connected to the endpoint AAS, e.g., a tablet.

The Java service environment also provides a `ServiceMapper`, a helper class that binds a service against a given AAS command server. Moreover, the `ServiceMapper` registers also the available metrics (see below) in the AAS command server. Ultimately, the `Starter` realizes a basic (optional) process to register all services given in a YAML service descriptor and to start the AAS command server on a given port. Services may also take care of a self-registration as it is the usual approach for Spring-based service implementations.

The right side of Figure 35 illustrates the extensible resource and service metrics framework based on the work of Miguel Gómez Casado [3]. All information to be monitored is represented in terms of gauges, counters or timers as defined in the micrometer monitoring interface⁸². We opted for micrometer, because a number of well-known monitoring tools such as Dynatrace support this set of concepts and because Spring Cloud Stream already exposes several default metrics via this interface. In more details, a *gauge* is a handle to get the current value of a monitored property, e.g., the number of threads in a running state. A *counter* is something that can be incremented or decremented by a fixed amount, while a *timer* is intended for measuring short-duration latencies, and the frequency of such events.

Micrometer provides interfaces and basic implementations for these concepts for the provider/service side, a JSON format to transport the information and through Spring a server to expose this information in terms of REST. However, the use of the micrometer ends at the service side, as typically accessing the monitored information is not part of the interface. Usually, the information is requested

⁸² <https://micrometer.io/docs/concepts>

through some form of REST client. This would contradict our basic requirement to try to realize all (distributed) communication via AAS (R7) or Industry 4.0 protocols (R14). Therefore, we use the REST information only locally and map this information into the AAS of the respective services. Moreover, it would be convenient to access the data of the meters in uniform manner also in other platform components, e.g., in the platform monitoring component. The first step towards this goal is to realize a request-side implementation of the micrometer representing (distributed) meters (package `environment.metricsProvider.meterRepresentation`). The second step is to encapsulate the communication, i.e., the micrometer REST communication as well as the AAS representation of the metrics. This is done in terms of `MetricsExtractorRestClient` and `MetricsAasConstructor` in `environment.metricsProvider.metricsAas`. Finally, the `MetricsProvider` in `environment.metricsProvider` defines the unified access to predefined micrometer elements such as the system memory, but also custom meters, e.g., to measure the stream throughput. Services can be explicitly marked as `MonitoredService` to receive an instance of the metrics provider in order to define and measure application-specific metrics. For example, the `AbstractProcessService` discussed above is a `MonitoredService` to provide access to the runtime metrics of the underlying process through OSHI⁵³. Moreover, services can be marked as `UpdatingMonitoredService` if regular updates of the measurements are needed. We consider monitored values as properties of the respective AAS submodel elements. To obtain the values, we attach functors to these properties to read out the monitored values. These functors may either rely on polling individual values via VAB [3] (pure R7) or on pushing the entire metrics provider via the Transport Layer (R7 and R14) into a local data instance attached to the AAS functors [4], where the performance impacts of three AAS integration patterns are discussed.

The classes discussed so far are intended for generic stream processors. Spring Cloud Stream and Spring Boot require specific code for services, the integration of the metrics and for their startup process. While the built-in metrics can be activated through a setting in the Spring application setup and by adding a respective dependency, the additional metric mechanisms defined in `environment.metricsProvider` are not automatically integrated. As discussed in more detail in [3], this is handled in the VAB poll approach by the extended `MetricsProvider` for Spring, the `RestAdvice` and the `MetricsProviderRestService`. Moreover, the startup code in `Starter` hooks into the Spring startup process, i.e., it obtains the Spring Rest server port, it attaches the port to the `MetricsExtractorRestClient` used by the upcoming services and starts the AAS command server of the parent class at a point in time when this is permissible for Spring. The `Starter` class is then integrated by the code generation in the actual service start code, which finally consists of just a few methods refining or delegating work to the `Starter` classes defined in the service environment.

3.7.3.2 The Python Service Environment

So far, we exclusively discussed the Java side of the service environment. Except for the monitoring and the Spring-specific implementation, the **Python service environment** is to a certain degree a mirror of the Java service environment. Differences are:

- The Python environment is accessed through the Java representation of services in the streaming engine, i.e., the Python environment realizes some form of command server as well as the soft-realtime data transport.
- A second difference is that we do not plan to monitor the non-Java environments unless explicitly required, because stream measures can be taken on the Java side. Resource measures such as memory consumption can be combined with the related Java process, i.e., the monitoring there requires an extension so that the resources consumed by the Python process can be considered. The Python monitoring is not part of this release.

As illustrated at the bottom of Figure 35, the Python service environment implements similar concepts as the Java environment. However, the Python service environment does not need to be a complete mirror as only the parts to execute services and to enable the communication between the Java side and the Python service implementation are required. Therefore, the Python service environment provides the `ServiceEnvironment` class, which imports service implementations dynamically from four packages named `datatypes`, `serializers`, `interfaces` and `services`. The first three packages are generated from the configuration model and provide datatype implementations, related serializer/type translator implementations and service interfaces specified in the configuration model (implemented based on `Service/AbstractService` from the service environment). The `services` package contains the manual implementations of the services.

For the data transport between the Java side and a service environment in a different programming language, e.g., Python, we can imagine the following alternatives:

1. Use of the platform Transport Component using a local transport server/broker. However, as the Transport Component is open regarding the transport protocol and the serialization, this would imply that all service environments (except for the Java service environment) must implement all transport protocol variants and all serialization mechanisms. If there is not just “the Python environment” but further language-specific environments, this leads to a plethora of required protocols. In particular, if the user organization decides to implement own protocols, these protocols must be mirrored into each environment. We do not think that this is a feasible and sustainable solution and opted for restricting the transport layer to Java code. Thus, the communication with Python may be based on a (local) protocol here, e.g., HTTP/REST as well as the serialization mechanism may even be fixed, e.g., JSON. While this approach is feasible as only a few variants are needed, it requires server processes for the communication on Java and Python side.
2. Extend the VAB protocol to transfer data. This is a specific decision of fixing protocol and serialization as mentioned in alternative 1. Also here, the back channel would require further server processes on the Java side. Moreover, as mentioned above in Section 3.6.2.3, VAB does not support for data transport and soft-realtime capabilities are questionable. We may take the VAB protocol into account. However, extending an external protocol also imposes compatibility and sustainability risks if decisions for the underlying implementation (BaSyx) are made, that conflict with our decisions.
3. As the non-Java service implementations are executed as local processes, also command line input- and output streams provided by the operating system may be a low-risk option (as some service candidates and many Unix command line programs do). Here, in particular the Java side must carefully parse the output of the executed services/service environment not to confuse “normal” output with data output. However, command line streams are said to be a performance issue on Windows-based systems.

Currently, the Python `ServiceEnvironment` implements the third alternative using (generated) object-to-string serializers with JSON as default wire format. We also use the command line streams for the command protocol. On the Java side, specific classes are bound against the Python service environment and the service deployment descriptor specifies the required service-specific Python artifacts as well as the command line parameters. More specifically, as already indicated above, the `PythonAsyncProcessService` is responsible for continuously running the Python `ServiceEnvironment` and the `PythonSyncProcessService` is an experimental call-and-return implementation of a Python service integration. While the `PythonAsyncProcessService` transmits data and commands to the `ServiceEnvironment`, the `PythonSyncProcessService` transfers only data items and calls Python upon each data item. Both integration classes support synchronous

3.7.3.3 Validation

So far, no performance evaluations of the generated code and the underlying service environment have been conducted. Therefore, the manually implemented service chains from the experiments discussed in Section 3.6.3.3 could be used as baseline.

```

stateDiagram-v2
    [*] --> Available
    Available --> Deploying
    Deploying --> Created
    Created --> Starting
    Starting --> Running
    Running --> Stopping
    Stopping --> Stopped
    Stopped --> Undeploying
    Undeploying --> [*]

    state "Normal Operations" {
        Running --> AdaptationPassivating["«Adaptation» Passivating"]
        AdaptationPassivating --> AdaptationReconfiguring["«Adaptation» Reconfiguring"]
        AdaptationReconfiguring --> AdaptationActivating["«Adaptation» Activating"]
        AdaptationActivating --> Stopped
        AdaptationReconfiguring --> AdaptationPassivating
        AdaptationReconfiguring --> AdaptationMigrating["«Adaptation» Migrating"]
        AdaptationMigrating --> AdaptationPassivated["«Adaptation» Passivated"]
        AdaptationPassivated --> AdaptationReconfiguring
    }

    Running -->|DeepHistory| Created
    AdaptationActivating --> Stopped
    AdaptationMigrating --> AdaptationPassivated
    AdaptationPassivated --> AdaptationReconfiguring
    AdaptationPassivating --> AdaptationReconfiguring
    AdaptationReconfiguring --> AdaptationActivating

    state "Failed" {
        [*] --> Recovering
        Recovering --> Recovered
        Recovered --> Recovering
    }

    Deploying --> Failed
    Stopped --> Failed
    Failed --> Deploying
    Failed --> Undeploying
  
```

Figure 36: Service states (comments cropped)

While the readings for the monitoring work fine, they are also just taken for individual measures rather than for full AAS. Polling all monitoring information from an AAS may induce a significant response time. Moreover, as detected during tests with the full platform components, a remotely deployed AAS reads out all values of all properties during its serialization for client use. This implies significant overhead and, more dramatically, in some cases even hang up the component. Thus, as indicated above, we realized a second approach based on turning the metrics provider into a micrometer-based JSON format and push this information into the (serialized remote) AAS. To avoid the problems mentioned before, the information is not written into the functors or the AAS rather it is implemented as a data instance shared by all functors of the service (represented by a submodel element). This shared instance can quickly react during AAS serialization (significantly faster than the 4-5 ms mentioned above) and initializes the transport connector lazily upon the first request. Moreover, this approach decouples the startup of the AAS implementation server as only the platform transport broker/server is used, which was already started along with the platform. Alternatively, for a plain AAS realization (R7), we could have realized a similar push approach through an operation on server side, but refrained from this idea as such a collector method shall not be part of the visible interfaces of the platform.

However, the push approach via the Transport Layer could leave the impression that the work on individual meters in an AAS is superfluous, in particular as the mechanism could equally be used to realize the central platform monitoring (cf. Section 3.8). This is not the case as discussions with other AAS users show: Nowadays, typical AAS tend to expose a huge amount of static and dynamic data for the described assets, i.e., the expected/actual resource consumption is often mandatory in some form.

3.7.4 Service Control and Management

The Service Control and Management component defines the service-interface of a (compute) resource towards the platform. It must provide means to load a service implementation onto the resource (in terms of binary artifacts, e.g., from a central platform server), to identify the descriptive information about services (id, name, description, version, service kind) and to provide access to runtime capabilities, e.g., the state of the service, reconfiguration capabilities, or runtime monitoring values. As the execution of the services happens within their (programming-language) specific environment, the control and management component can be realized in generic manner.

Individual services must comply with a lifecycle that can be queried and influenced by the platform. The underlying lifecycle state machine is depicted in Figure 36. Services can be downloaded from the service store and become **Available** on the hosting resource. When triggered through the platform and the **ServiceManager**, a service is deployed (**Deploying**) and gradually turns into the **Running** state. If nothing bad happens at runtime, a service is stopped through the **ServiceManager** (turns to **Stopping** and **Stopped**) and if requested, may be removed from the resource (**Undeploying**, afterwards **Unknown**, not shown in Figure 36). At runtime, a service may be reconfigured, adapted or migrated (which may need passivation and activation). Further, a service may fail, which can lead to a recovery procedure (in the lower sub state machine in Figure 36). If the service becomes operational again, it jumps back into the upper sub state machine and there into the last “normal operation” state (via the UML H* deep history state) and goes on from there.

Figure 37 illustrates the design of the Service Control and Management component. At the core of this layer is the **ServiceManager**, which performs operations such as starting and stopping individual services. While the transitions displayed in black in Figure 36 are controlled by the **ServiceManager**, the transitions in red are performed by the service and monitored (via AAS property value polls) by the **ServiceManager**.

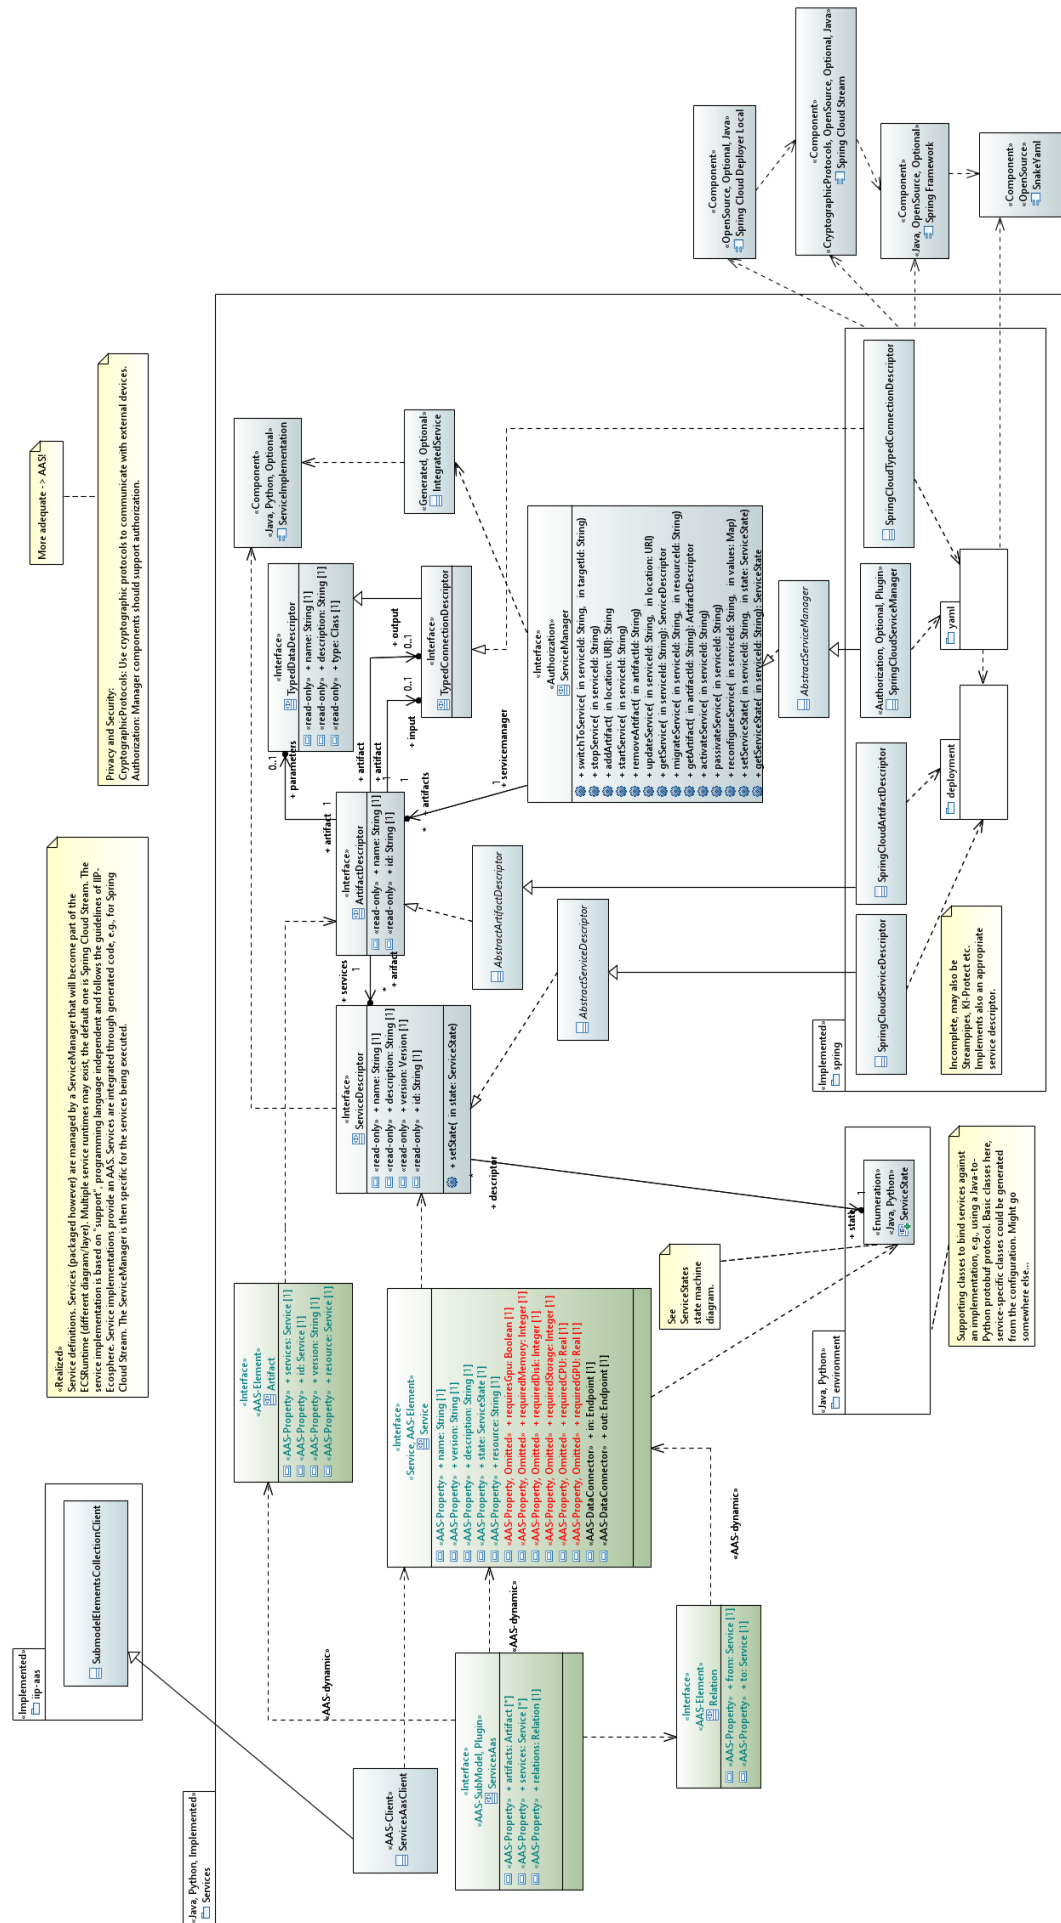


Figure 37: Service interfaces and management

Services are packaged and transported in terms of artifacts, i.e., an artifact may contain multiple services realized in different programming languages. Instead of the actual instances that may be located in a different container, the ServiceManager primarily operates on descriptors, such as the ArtifactDescriptor detailing structural information on contained services. Access to artifacts and services happens through identifiers, whereby several operations and information accesses are delegated by the service manager to the descriptors or through the respective AAS to the AAS implementation server directly approaching the respective service instance.

Figure 38 illustrates the structure of a generated application artifact for execution with Spring Cloud Stream. Here, an artifact consists of combined binaries provided by the service execution framework (the startup code) and binaries that make up the application (below BOOT-INF, requiring a specialized ResourceResolver, cf. Section 3.5.6). Within the application parts, such an artifact contains all application dependencies (in lib) including, e.g., the service manager of the IIP-Ecosphere platform and its transitive dependencies, but also generated parts such as the application interfaces (folder iip, also containing the startup class for the Java service environment Starter.class), the Spring Cloud Stream application specification (application.yml), the service deployment descriptor (deployment.yml) and the logging configuration (logback.xml). Depending on the integrated services, more artifacts may be included, e.g., reusable service binaries (here kodex.zip), customized service artifacts (kodex_pseudonymizer.zip) or the Python code for executing a Python service in the Python service environment (python_kodexPythonService.zip). Such specific binaries are referenced in the deployment descriptor and unpacked by the service manager upon start.

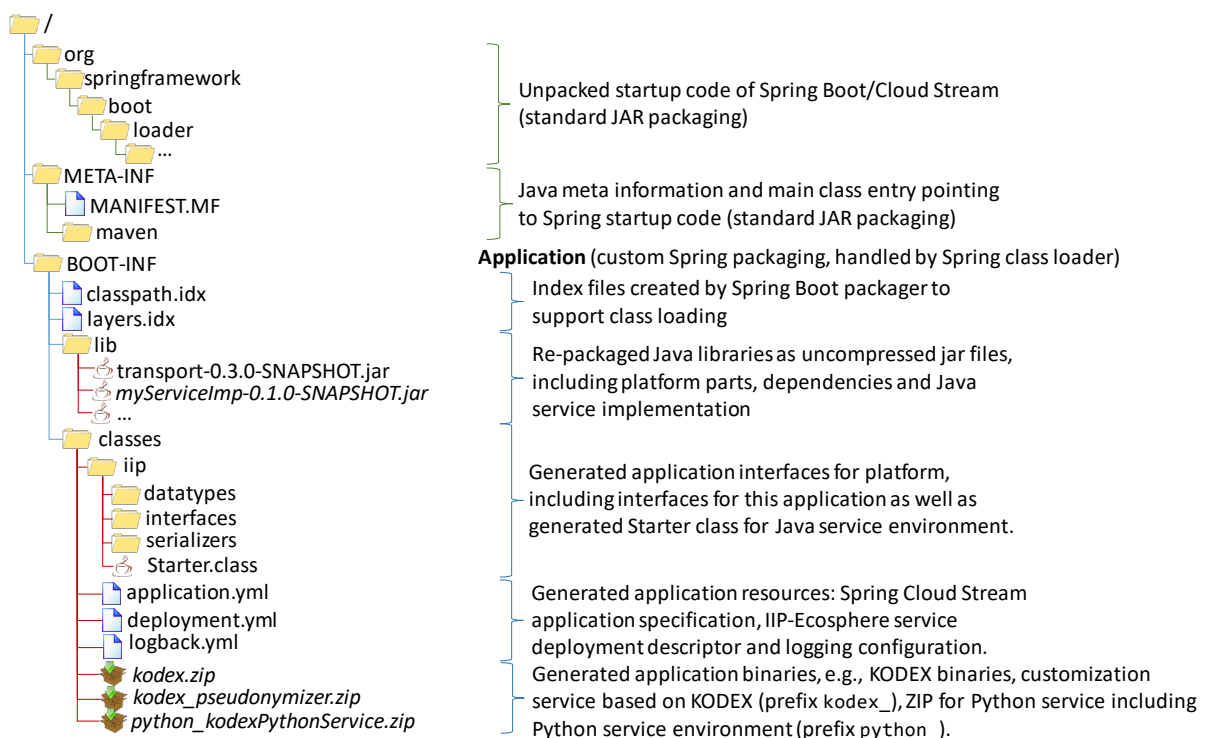


Figure 38: Structure of a JAR application artefact for the Spring Cloud Stream engine.

The packaging shown in Figure 38 represents an executable Java ARchive (JAR) and may even be executed without the platform (provided that a respective setup, e.g., communication ports are given). In principle, this is also one major functionality of the service management - besides passing environment settings such as (dynamic) ports to the services stored in the service artifact. However, as an executable JAR, the packaged Spring application sets up an own class loading mechanism, which prevents shared libraries, i.e., libraries that must not be packaged into the artifact to reduce the footprint of the artifact. However, modifications to this startup procedure are not permitted, e.g., an

additional, earlier class loader. To allow for shared artifacts, the service management supports a secondary format, which is independent of Spring. This packaging structure is shown in Figure 39.

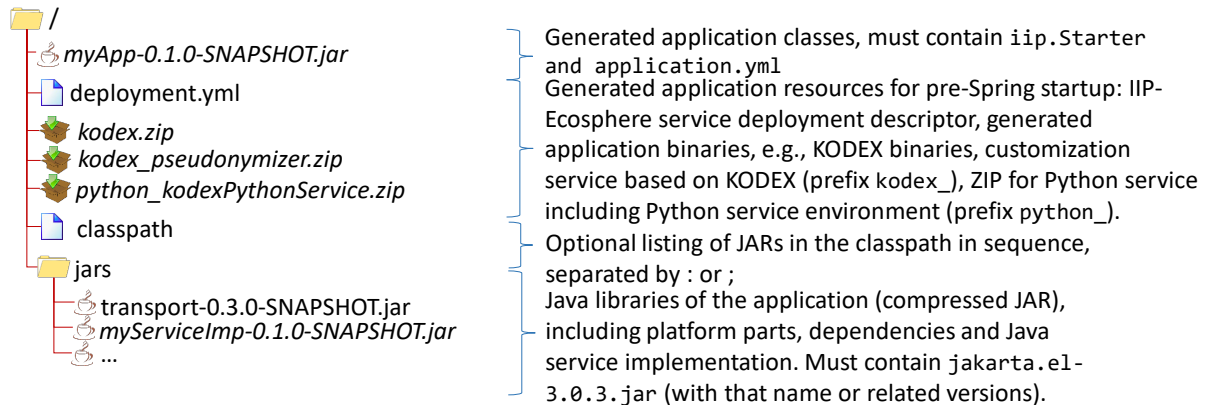


Figure 39: Structure of a ZIP application artefact allowing for shared libraries (as variant of Figure 38).

In contrast to Figure 38, the ZIP-based artifact is not an executable JAR and contains packaged rather than unpacked JARs. The application JAR including the generated class `iip.Starter` must be in the top-level directory of the ZIP. There must also be the service deployment descriptor (`deployment.yml`) so that the service manager can read the information about the contained services. Moreover, also the “binary” artifacts, e.g., for KODEX or python must be on top-level, so that the service manager can extract the artifacts for executing them in terms of operating system processes. In contrast, the Spring application definition (`application.yml`) and the logging configuration shall reside in the generated application artifacts as they will be loaded by the respective Java libraries on demand via class loading. The dependencies of the application are located in the `jars` folder (or optionally on top-level). Here the difference to Figure 38 is that any shared jar can easily be removed from that folder (during the packaging process or manually for experiments) and provided through a shared libraries folder⁸³ known to the service manager. Optionally, the file may contain in the file `classpath` a listing⁸⁴ of Java libraries in their intended sequence to avoid conflicts.

The `ServicesAasClient` provides access to the properties and operations of the AAS of the service layer. To avoid adding even more visual complexity to Figure 37, we did not indicate the relation between `ServiceManager` and `ServicesAasClient`. Actually, both implement the same interface called `ServiceOperations`, which contains the basic operations of `ServiceManager` not requiring the (repeated, potentially inconsistent instantiation of) service descriptors. The `ServicesAasClient` can be used by upstream layers to conveniently access the services AAS.

Typically, services can operate as a single process and do not require further resources, e.g., a central server process to communicate with. However, there are services with this requirement, e.g., the anonymizer and pseudonymizer KODEX (cf. Section 3.9.1), which optionally may cooperate with a server instance, or a federated learning approach, which usually requires a central model management and exchange server. There are settings, where such server processes may be installed and controlled by the user, but we can also have application-specific server processes, which shall be integrated into the lifecycle of an application, and, thus, be started and stopped along with the application by the

⁸³ Currently, this folder is specified in the setup information of the service manager. This information could be relocated into the service deployment descriptor in future versions of the platform.

⁸⁴ Relative file names in Windows or Linux notation, separated by `:` or `;` depending on the operating system. Such a file can be created by Maven. Before execution, the file is rewritten to comply with the `@` argument file format of Java 9 and newer. Thereby, the JARs in the root folder of the archive are added by the service manager to the start of the classpath. If no such file is present, a wildcard classpath is constructed, which may cause accidental class loading conflicts.

platform. Moreover, certain applications may require some form of interaction with the server, e.g., to re-configure KODEX at runtime or to request the replay of a model snapshot of a federated learning server. It is desirable that such interactions happen in a uniform manner so that application components can rely on the approach. For the IIP-Ecosphere platform, it would, thus, be desirable to represent such server processes within the platform AAS and to allow, e.g., for runtime re-configuration of service parameters (which may also trigger a model exchange). Furthermore, it would be desirable that the resource which hosts the server process can, as for services, be decided freely within the available resources of a platform instance. Moreover, a uniform, standard-compliant communication between client and server would be desirable, most preferably via the Transport Component (cf. Section 3.6.2) of the IIP-Ecosphere platform to also exploit the flexible exchange of protocols. However, directing an externally determined communication via the Transport Component may require changes to the server/service to be integrated, which may not be possible in all cases, for which we prioritize this requirement as optional, i.e., to be decided based on the server/service at hands.

While we believe that in most cases the open service concept of IIP-Ecosphere is also applicable to services in a server/service setup, the server side requires specialized capabilities. In general, we consider a server as a realization of an (internal) service, which does not exchange data via the service input/output interfaces. This allows for utilizing AAS publishing, distribution, re-configuration and execution mechanisms of usual services, in particular the service environments for Java and Python, while not requiring an input/output modeling for the internal service/server communication in the configuration model (cf. Section 6). For this purpose, some server-specific capabilities had to be added to the service environments, e.g., that services know the (optional) service/server communication channel they are relying on. It further requires specialized capabilities such as assigning server processes to an application (in the configuration model), registering and announcing the actual server network information via the platform network management (cf. Section 3.5.2), starting/stopping server processes along with applications (while maintaining the number of using service instances using these in the platform network management) or provisioning of specialized transport streams for (private) service/server communication. The required capabilities are partially realized in the service manager and partially in the service environment as generic or specific functionality of basic service functionality.

Different technologies can be used to realize and execute service chains, i.e., to efficiently pass data along pre-defined data paths between the services, to transform data where needed etc. As part of such a service chain, data is turned into some form that can be transported by the utilized protocols. This serialization as well as the transformation of data to fit the input/output requirements of a service is part of the mechanisms of the Transport Layer. As discussed in Section 3.6.2, we rely on Spring Cloud Stream as default (stream-based) service execution engine. An integration of the transport level protocols and serialization mechanisms for Spring Cloud Stream was introduced in Section 3.6.2. As also stated there, we foresee that the IIP-Ecosphere platform shall support also other service engines in a flexible manner. Thus, the design of the Service Control and Management Component must allow for the execution of the management binding against alternative service execution technologies. For this purpose, the `ServiceManager` as well as the related descriptors are defined in Figure 37 as interfaces (in the package `services`), while the actual implementation is realized in a separate, alternative package (`services.spring` for the default engine), which hooks itself as implementation into the basic service management interfaces via JSL.

The default implementation of the `ServiceManager` in `services.spring` relies on Spring Cloud Stream, the Spring deployer mechanism and, in turn, on the Spring Boot framework. For Spring-based services, the packaging happens in terms of specifically packed JAR files (a form of “fat”, standalone JAR files also containing the required Spring base classes). Besides the Spring application configuration

defining the actual wiring of the services, such artifact JAR files contain a deployment specification detailing the services, their communication, service dependencies and, if required, also non-Java service implementations and their integration. Following the conventions of the platform, these deployment specifications are stated in terms of YAML files. Both, the Java object representation of the YAML contents as well as the JAR artifacts are linked against the Spring-specific Artifact and Service Descriptors, which contain additional information required to manage services using Spring. The Spring-based `ServiceManager` utilizes the Spring deployer mechanism, i.e., the local Spring deployer. The deployment specification also allows defining external service implementation processes, e.g., for Python, so that the data communication is managed by Spring-services while the actual implementation of the service operates in an own process. By default, services are executed in their own processes so that services can be restarted in case of failures (R9) without accidentally shutting down healthy services. However, such a single-process deployment may not be desired in some cases so that the deployment descriptor allows for specifying “ensemble” services, i.e., Spring-services that must be executed within the same process.

The AAS for the Service Layer consists of a services sub-model indicating as sub-model element collections the (locally) installed artifacts and the contained services, the installed services and their properties as well as the data paths/relations among services. When a service is started, its state changes and for each data path a relation instance is created, i.e., a relation represents the instantiated data path between two service instances and points to the actual start and end service. Start and end service occur in the AAS as soon as the respective service is created. In turn, this information is used by the service manager to determine available services, e.g., during startup of dependent services in service chains. Most operations provided by the `ServiceManager` (also via AAS) are parameterized by an artifact or service identifier. However, internally the operations are bound to the resource the respective artifact/service is installed on, so these operations do not occur at the services in the services collection rather than for the resource in the resource collection. We will detail the resources in Section 3.8.1 as part of the design of the ECS runtime. As all those operations may fail, the implementation must not only return a result but also carry information about thrown exceptions when calling an AAS operations.

The service manager AAS is primarily intended as service-level control and monitoring interface. Services are supposed to register themselves with the respective local AAS command server (see also Figure 5 in Section 3.1) to react on command requests. Similarly, when monitoring information is requested, the (central or locally deployed) AAS communicates with the respective AAS command server. In case of services not implemented in Java, the respective service environment must provide an AAS command server and pass the information on to the service instances.

As discussed above, soft-realtime data streams shall not be transmitted through AAS rather than through the streaming engine (for our default engine using one of the protocols of the transport layer). If the service implementation is done in Java, the streaming engine will directly communicate with the service (potentially involving glue code generated from the platform configuration). If non-Java service implementations are used, the service representation in the streaming engine must route the data to the respective service environment, which shields the services from the actual communication and passes the data in adequate form to the respective service instances.

The requirements in [13] do not explicitly define the properties that shall be monitored for services. R29a, R70, R122f just indicate that services may have quality properties, e.g., to support adaptive service selection. Monitoring probes may be generic or bound to the services and, thus, are realized in the service environment (in particular the default one for Java, cf. Section 3.7.3). Similarly, the creation of related parts of the AAS are realized there. Further, probe services may be inserted to perform application-specific monitoring. However, probe services are currently not realized.

The ServiceManager and its AAS are validated in terms of regression tests. As the ServiceManager and the descriptors are interfaces/abstract classes, the validation must set up a pseudo implementation for basic testing. The Spring Cloud Specific functionality is tested through a handcrafted service artifact with simple contained services and multiple deployment descriptor targeting different artifacts, e.g., with or without process ensembles. This artifact is based on the Java service environment (cf. Section 3.7.3). In these tests, the setup of the ServiceManager provides a broker, dynamic network settings are handled by a local NetworkManager and the service manager is utilized for starting and stopping services. The running services are validated in terms of their data throughput and the actual metric values that the services provide, i.e., that the metrics defined in the service environment (cf. Section 3.7.3) become part of the AAS of the service management. Moreover, also the dynamic aspects of the AAS are validated, in particular during startup in order to figure out whether a service is already running.

Furthermore, the Spring Cloud based Service Manager was validated in a Linux VM-based server setting as well as on a Phoenix Contact AXC 3152⁸⁵ PLC-Edge with 2 GByte RAM and 8 GByte memory card providing additional hard disk space. On the Linux server, the Service Manager was executed directly on the operating system as well as in a Docker Container, on the AXC we focused only on the Docker setup. In both cases, we were able to manage a simple demonstration application (adding the artifact, starting, stopping, removing the artifact) and to verify that the expected input/output behavior of the services can be observed. As starting and stopping individual services involves powering up a JVM, the service manager takes a certain operation timeout (with a default length of 30 seconds) into account. This is sufficient for the Linux server (and the regression tests mentioned above), but does not work on the AXC 3152, where a longer timeout is needed.

Currently, as prescribed by the development streams in Section 3.2, design and realization focus on the distributed management of services. Advanced capabilities such as switching among alternative services or migrating services are subject to a later release.

⁸⁵ <https://www.phoenixcontact.com/online/portal/de?uri=pxc-oc-itemdetail:pid=1069208&library=dede&tab=1> provided by Phoenix Contact to the SSE group of the University of Hildesheim <https://sse.uni-hildesheim.de/aktuelles/detailansicht/weltweiter-marktfuehrer-unterstuetzt-universitaet-hildesheim-im-bereich-industrie-40/>

3.7.5 Requirements Discussion

In this section, we review the already realized requirements for this layer. As mentioned in the sections before, we aimed for a basic implementation in this release, which is reflected accordingly in Table 11.

Table 11: Review of realized⁸⁰ requirements for the Service Layer (based on Table 2, Table 3 and Table 10)

Requirement	Summary
R4a	Services with input/output in ServiceDescriptor and AAS (services sub-model)
R4b, R131b	Metadata in ServiceDescriptor and AAS (services sub-model)
R4c	State in ServiceDescriptor and AAS (services sub-model)
R4d	The generic execution state of Spring Cloud Stream is supervised, also the resource consumption and stream processing or applications can define custom metrics. The state is reflected into the service descriptor, state and monitoring information (as far as enabled) show up in the AAS (services sub-model). <i>So far, services are not re-started if they fail as this is not provided by the Spring Local Deployer.</i>
R4e	<i>Not yet implemented, but possible e.g., through service descriptor.</i>
R4f	Possible via probe services with specific interface.
R20	Application-specific data paths are supported through the streaming library (by Spring Cloud Stream, even at runtime). Data paths are dynamically indicated in the AAS (sub-model relations).
R20b	Data paths can have properties/parameters. Basic properties like the protocol or the encoding are supported by Spring Cloud Stream. Additional properties can be specified in the service deployment descriptor (Yaml file in artifact).
R20c	Data paths shall be managed by the platform through the configuration (Section 6), code generation, Spring application setup/service deployment descriptor and during service startup, e.g., taking dependent services into account.
R29a, R70, R122f	Services describe their functionality and their runtime properties (as provided, selected, implemented) through the Service Management and Control AAS, in particular supported by service monitoring.
R31	A container shall contain only the required components/services. This is supported through the service artifacts/containers, that are generated from the configuration model for a certain target deployment, i.e., with the (minimal) required resources.
R31	Required components and dependencies can be specified in the configuration
R31b	Artifacts may contain optional components, which are then not executed. <i>Optional services and their wiring is not subject of this release.</i>
R31c, R69	Dynamic exchange of service implementations is prepared by separating service implementation and (generated) binding against the stream processing library. The service interfaces allow for dynamic exchange and service migration and the operations are available through the AAS of the service management and control component. <i>The realization of the operations was not part of this development stream/release.</i>
R32	Services can declare and describe typed parameters. The ServiceManager supports changing these configuration parameters.
R38-R68	<i>A variety of security and privacy mechanisms are introduced (e.g. in Figure 20) which ensure relevant security and privacy requirements. Security mechanisms are indicated in the architecture but not part of the platform implementation of this release.</i>
R73a-f, R79	As argued for the Transport/Connector components, we do not limit or prescribe types.
R113	In the first place, the ServiceManagement and the stream processing approach are both realized in Java. Services realized in Java can be directly integrated with the stream processing and be executed in the same threads/processes. Python

Requirement	Summary
	Service implementations can be integrated through generic processes, classes and the Python Service environment, so far based on command-line process streams. The use of specific Python frameworks as requested in R113a belongs to the realization of the AI service toolkit.
R132	Application-specific services are supported through service interfaces as well as integration of artifacts in the code generation and packaging process.
R133	Runtime support for applications (status of services) is provided via the ServiceManager and the AAS.
R133a	The status of services is provided via the ServiceManager and the AAS
R133c	Within the limits of the service state machine, the ServiceManager and the AAS provide means for adjusting the state of a service. In particular, functions for activating, passivating and migrating services are provided and generically implemented.
R133d	<i>No specific functionality to resolve error conditions is provided in this release.</i>
R134c	The ServiceManager supports stopping as well as removal/disposal of services and service artifacts.
R135	An operation for updating services is provided and available through the AAS, <i>but the operation itself is not implemented in this release.</i>

We conclude, that most of the basic requirements for this layer are already implemented. Advanced functionality such as dynamic service operations or monitoring as well as security and data protection mechanisms are subject to one of the next releases.

3.8 Resources and Monitoring Layer

The Resources and Monitoring layer enables the deployment of services to (edge, server, cloud) devices, allows for overarching management of the devices and provides aggregated monitoring information about running resources and services. Moreover, the first platform components for the overall management of resources, namely the device management and the platform monitoring are located in this layer. We will discuss the ECS runtime in Section 3.8.1, the device management in Section 3.8.2 and the monitoring in Section 3.8.3.

3.8.1 ECS runtime

Flexible and heterogeneous deployment to edge, server and cloud devices is a central capability of the IIP-Ecosphere platform. [13] defines several requirements for the envisioned deployment approach. Table 12 summarizes the requirements to be considered. R25c and R25d target the (central) management of resources and, thus, are addressed by the device management in Section 3.8.2.

Table 12: Specific requirements for the heterogeneous deployment (excluding configuration)

Requirement	Summary
R23	Support for dynamic deployment
R24	Deployment to different types of resources/hardware
R25	Resource properties or functionalities described as AAS
R25a	AAS of a resource shall be realized by an ECS runtime
R25b	AAS of available resources must be announced to the platform
R25e	Resource AAS must describe static properties
R25f, R103a	Resource AAS must describe dynamic properties, e.g., the resource utilization, the memory usage or the utilization of CPU/GPU/TPU cores
R25g	Resource AAS must contain functions for the deployment
R25h	Resource AAS shall contain functions for exchanging deployment units at runtime
R26	Deployment to on-premise resources

Requirement	Summary
R27	Optional deployment to connected IIP-Ecosphere instances
R28	Optional deployment to cloud resources, e.g., Google Cloud or Gaia-X
R29	Deployment unit must provide an explicit interface in terms of an AAS
R29a	(Quality) properties and functional interfaces of deployment unit via AAS
R29b	Deployment unit AAS shall be linked to resource AAS
R29c	Contained services/containers shall be made available via the deployment unit AAS
R30	Deployment unit must be encapsulated as container
R30a	Deployment units on IT level must be technologically uniform
R30b	Deployment units on OT level can be technologically different
R30c	Platform can support the integration of external container repositories
R31	Container shall contain only the required components/services
R31b	Containers can contain optional components
R31c	Components/services in a container may be exchanged dynamically
R32	Container can contain data/models, to be configured via parameters
R33	Container can contain local data stores
R35	Sampling rate of 2 ms through container
R36	Optional configuration of resources
R36a	Writing of resource configuration
R36b	Reading of resource configuration
R37	Optional remote maintenance of resources
R38-R44	Security requirements
R45-R68	Data protection requirements
R102	The platform must allocate deployment units to available resources.
R105	Deployment must happen via platform functions and AAS.

As described in [41], each device shall execute a basic runtime component (ECSRuntime) providing the AAS of the device and managing the containers in which individual services are executed. Figure 40 illustrates the design. The fundamental parts are the ResourceUnit representing the AAS of the resource on which the runtime component is executed as well as the DeploymentUnit containing the services executed on the resource. The Service Management and Control component from Section 3.7.3 contributes information to the DeploymentUnit, e.g., the running services and their instantiated relations. Through the ECS runtime, the device can receive and execute commands from the platform, such as downloading or starting a container. Moreover, different container technologies must be considered and addressed in a uniform manner through the ECS runtime.

Different ways to install such an ECS runtime are possible depending on the capabilities of the underlying device. The default approach is to provide an automatically created container with the instantiated ECS runtime as well as one or multiple (dynamic) containers for the services. Depending on the capabilities of the device, e.g., whether a suitable version of Java is available, the ECS runtime could also directly be installed on a device. If in the future a resource description such as the AAS of the ECS runtime is standardized (and IIP-Ecosphere platform is compliant with that standard), one could also imagine that the device already ships with an ECS runtime (possibly realized in some other language than Java) or it can be installed from the store of the device vendor. Measures to install, manage and update such installations are subject to the device management (Section 3.8.2).

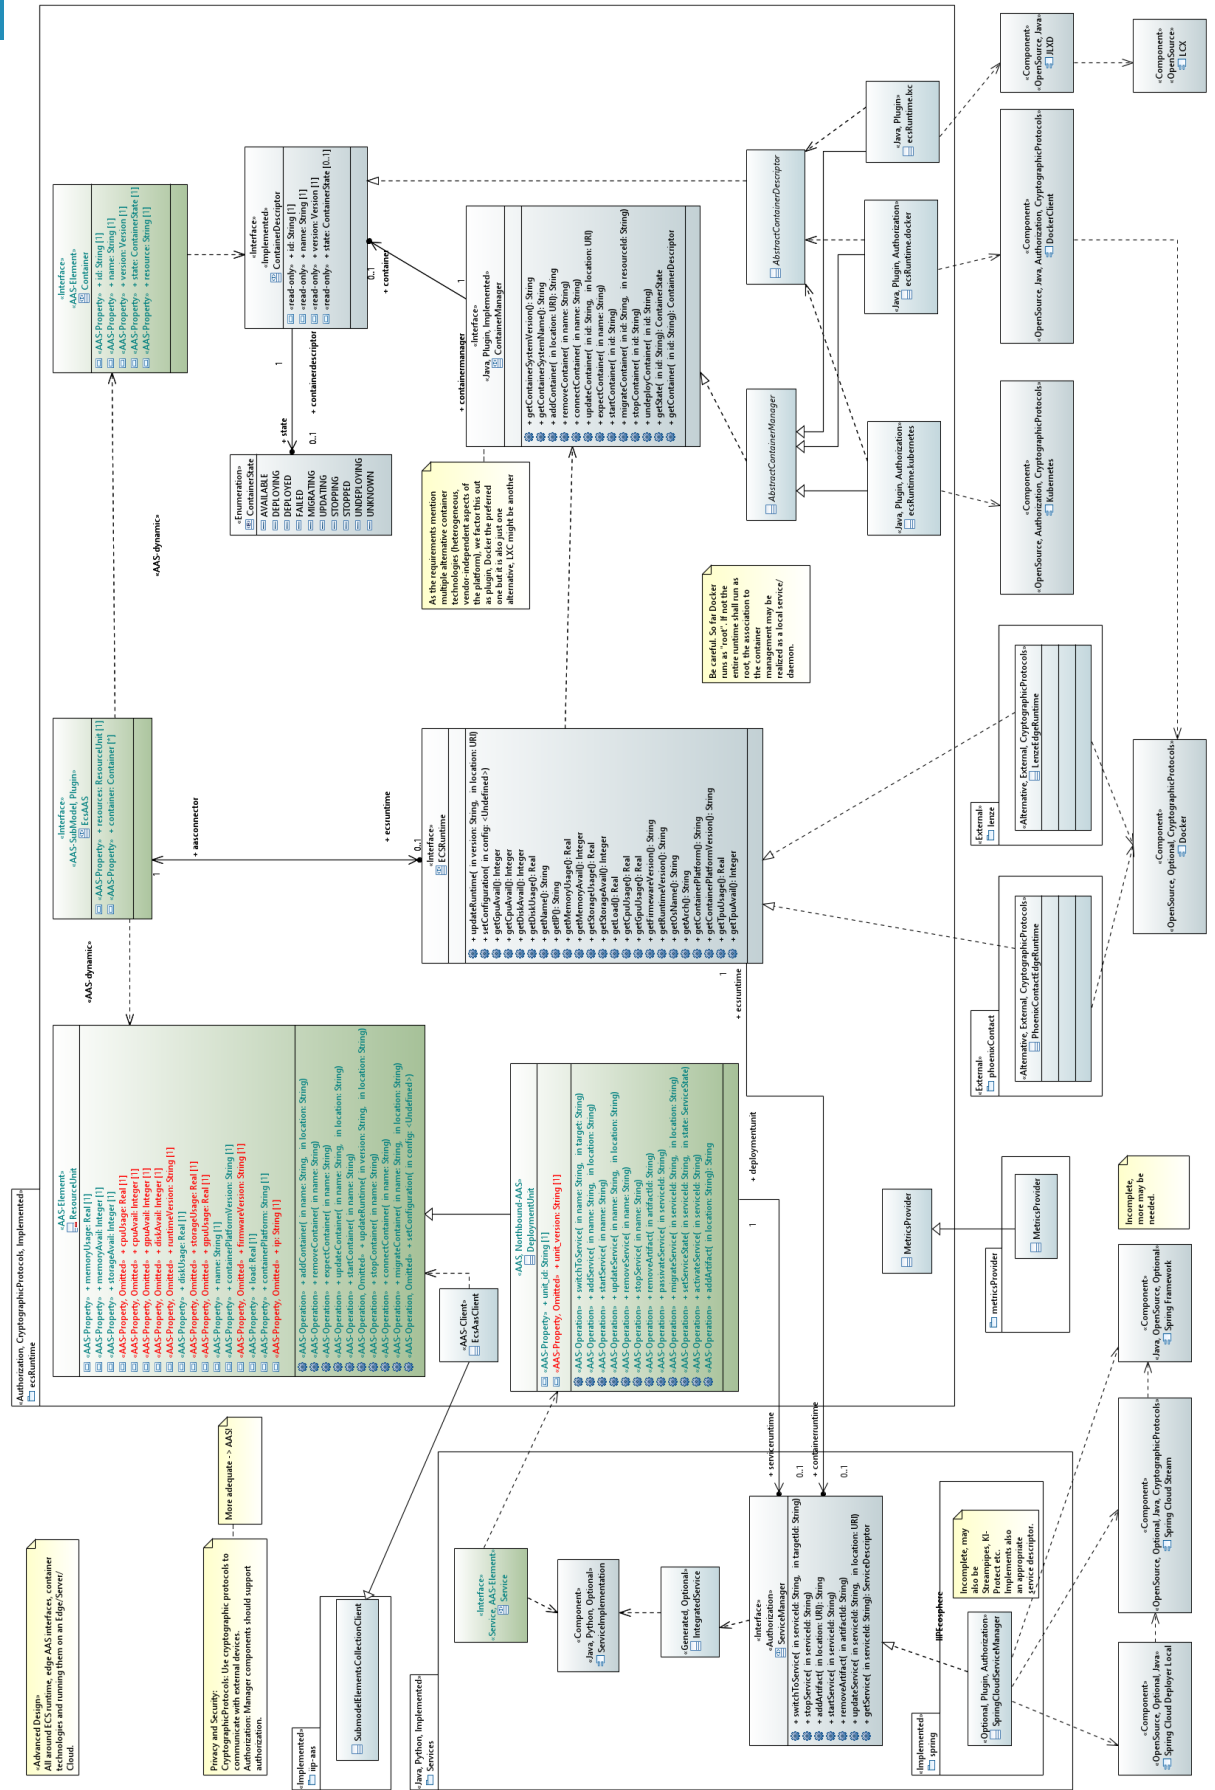


Figure 40: ECS runtime for Service Deployment (comments partially cropped)

As already emphasized in [41, 13], one fundamental basic work for the resource abstraction runtime is the LNI 4.0 edge configuration usage view [26]. [41] subsumes and extends [26] and [13] integrates relevant requirements from [41]. As the need for managing resources and containers on resources, in

particular edge devices, is known in Industry 4.0, platforms [35] and also other work address this topic in various ways. In addition to the 21 platforms analyzed in [35], also approaches like OpenHorizon⁸⁶, the IBM Edge Application Manager⁸⁷ or the ICP4Life platform [27] have been researched or are available. In recent time, also container orchestrators such as Kubernetes⁸⁸ became popular. Although there are significant overlaps, there are also important differences between these approaches and the ECS runtime in IIP-Ecosphere. One main difference is the general requirement R7 that all interfaces in the IIP-Ecosphere platform shall be based on AAS aiming at an interoperable integration of heterogeneous devices (based on an agreed structure, at least within IIP-Ecosphere). Moreover, it is important to point out that IIP-Ecosphere aims at a flexible integration of components leaving the final decision to the installing company through the configuration model, i.e., we do not make decisions such as statically relying on Kubernetes. In contrast to existing edge management approaches, as already pointed out in [41], IIP-Ecosphere aims at supporting more sophisticated management operations on the edge, in particular for data paths and relations as discussed in Section 3.7. Also platforms have been created as part of research projects, e.g., ICP4Life [27]. However, not all of these platforms are publicly available and suffer from similar overlaps and differences, e.g., the strong focus of IIP-Ecosphere on AAS and AI.

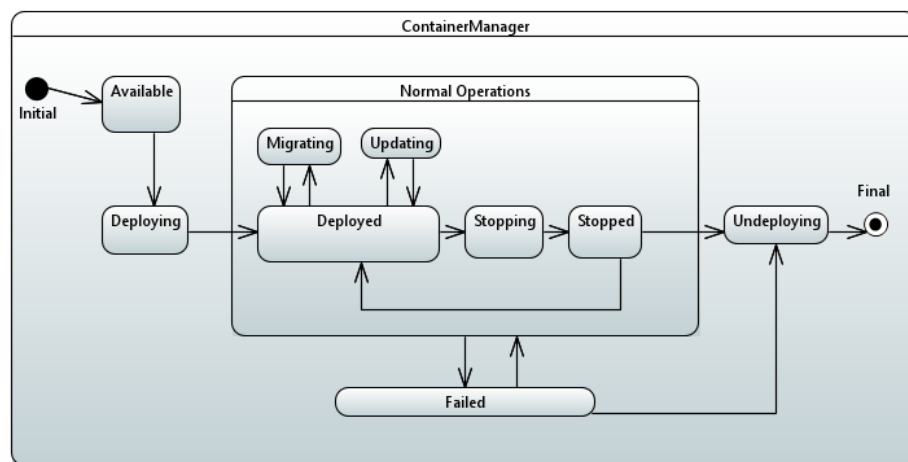


Figure 41: Container states

Figure 40 illustrates the design of the ECS runtime component. Figure 3 in Chapter 3 already discussed the context/stack for this component, i.e., on top of the AAS support, network management, transport and connectors and (optionally) service layer, the ECS runtime is supposed to provide a resource abstraction to manage the containers containing services to be executed on a resource. At the heart of the component is the ECSRuntime which acts as internal façade³⁹ for the ECS runtime AAS. Behind that façade, the actual operations are realized to be able to customize the ECS runtime for the resource at hands. Two example devices (produced by Phoenix Contact or Lenze) are indicated in Figure 40, but also a GenericJavaRuntime, which relies on an abstract ContainerManager (along with a ContainerDescriptor, akin to the service descriptors in Figure 37). Akin to the service manager, the container manager provides AAS operations, e.g., to download, start or stop a container. The respective container states are depicted in Figure 41.

⁸⁶ <https://www.lfedge.org/projects/openhorizon/>

⁸⁷ <https://www.ibm.com/docs/en/edge-computing/4.1>

⁸⁸ <https://kubernetes.io/de/>

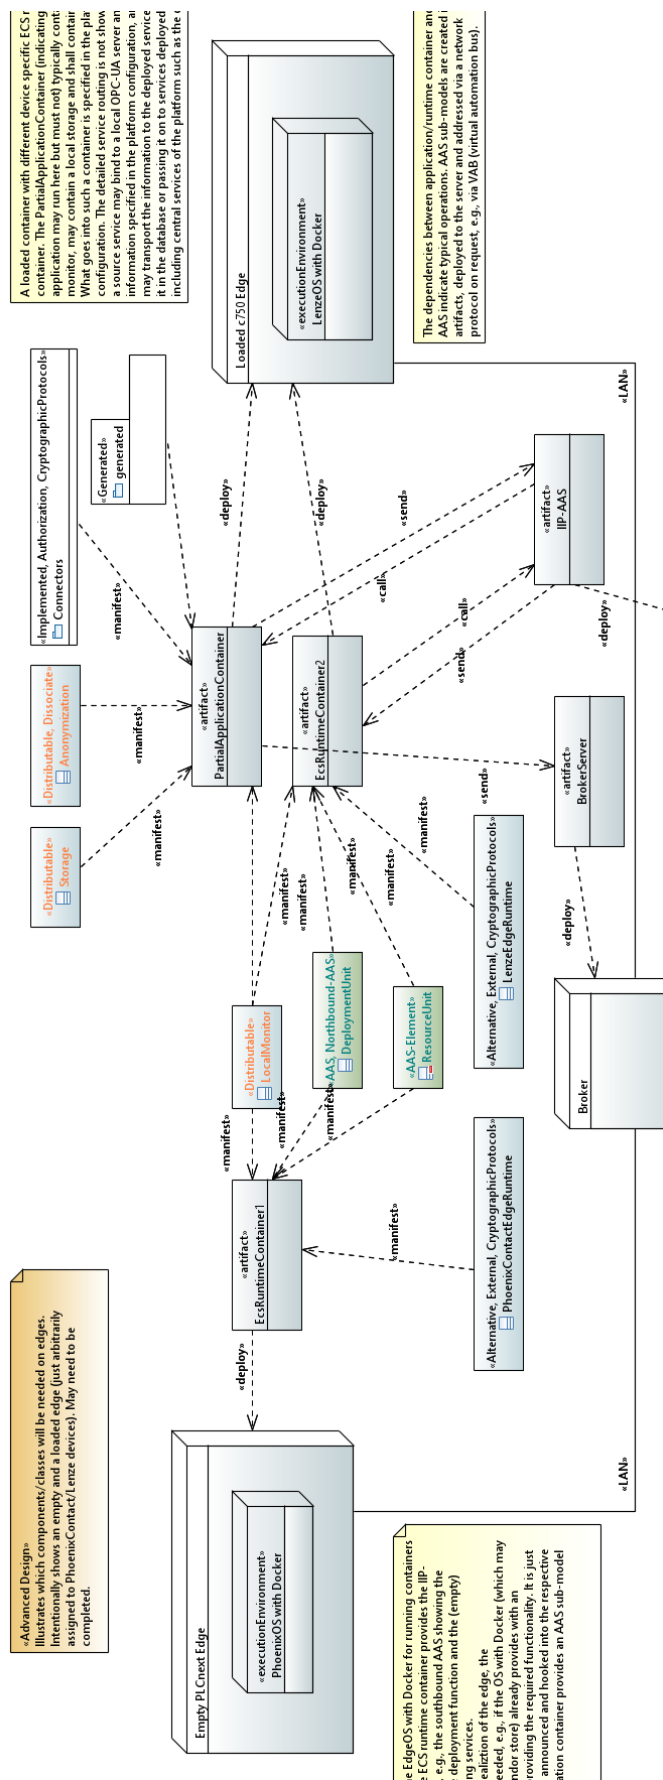


Figure 42: Example deployments to empty (left) and loaded (right) edge device (comments and further deployment nodes representing the platform partially cropped)

Partially, device-specific functionality is supposed to be realized via other plugins, e.g., target-system specific implementation of the `SystemMetrics` (cf. Section 3.5.4) or how to provide and access to the device AAS. It may be that the device is already older and does not provide an AAS. For this purpose, the ECS runtime allows to customize the AAS origin via the `DeviceAasProviderDescriptor` (not shown in Figure 40), which determines the component that returns the address of the respective device AAS (the component may also create the AAS if needed). Currently, three implementations of the related `DeviceAasProvider` are shipped with the platform, a pragmatic one reading the AAS from a simple Yaml file (reading manufacturer/product images from the same location), an implementation reading the AAS from an AASX package file and a multi-provider that selects the first provider (by default Yaml or AASX) that returns an AAS address. For the Yaml/AASX providers, the underlying information is retrieved as classpath resource, either as `nameplate.yml/deviceId.yml` or `device.aasx/deviceId.aasx`, respectively whereby `deviceId` is taken from the `IdProvider` of the support layer. This is currently the most easy and compliant way to specify device-supplied services and their properties, e.g., port or protocol version.

The IIP-Ecosphere platform provides a plain Docker⁸⁹ container manager (`DockerContainerManagement`). As for the service descriptors, the `ContainerDescriptor` is manifested in terms of a Yaml file, which is supposed to form the main entry point for adding a container at runtime, i.e., the platform specifies a URI pointing to the Yaml file, which indicates the name of the packaged container at the same location. We refrained from adding the descriptor to the packaged container as this may not be permissible for some container formats. In addition, the IIP-Ecosphere platform offers an LXC container manager. We selected LXC [38] due to the use in other industrial platforms [35] and as Docker changes its (commercial) licensing. It is important to mention that LXC is licensed under GPL and, thus, must be an optional component of the IIP-Ecosphere platform due to platform licensing rules. LXC is integrated via the Java LXD library `JLXD`⁹⁰ and requires specific installation steps, which are detailed in [38]. In the current state, it is not supporting all container creation strategies and may be limited regarding a container registry. A container manager for Kubernetes via Industry 4.0 protocols (R7, R14a) is in development but not part of this release.

The `EcsAasClient` provides access to the properties and operations of the AAS of the resources layer. To avoid adding even more visual complexity to Figure 40, we did not indicate the relation between `ContainerManager` and `EcsAasClient`. Both classes implement the same interface called `ContainerOperations`, which contains the basic operations of `ContainerManager` not requiring the (repeated, potentially inconsistent instantiation of) container descriptors. The `EcsAasClient` can be used by upstream layers to conveniently access the ECS runtime AAS.

At a glance, Figure 40 does not indicate much monitoring support for the ECS runtime except for some AAS properties in `ResourceUnit`. As the Java service environment (cf. Section 3.7.3) provides a generic and extensible monitoring approach, we re-use it here although the ECS runtime is not a “service”. Thus, the ECS runtime defines a `MonitoringProvider` as well as a regular monitoring update operation that is started as part of the JSL lifecycle descriptor of the ECS runtime (not detailed in Figure 40). The operations to create the AAS refer to the `MetricsAasConstructor` of the Java service runtime mirroring a default set of meters of the monitoring provider into the AAS of the ECS runtime (therefore, currently some runtime properties in `ResourceUnit` are realized while others appear as omitted). Depending on future decisions, a specific set of resource meters can be defined and applied to both components in uniform manner.

⁸⁹ <https://www.docker.com/>

⁹⁰ The official repository is at <https://github.com/digitalspider/jlxd>. However, due to required bugfixes and the need for a deployment to Maven central, which was not provided by the original authors, we rely on a fork of `JLXD`, which is part of `EASy-Producer` <https://github.com/SSEHUB/EASyProducer>.

Figure 42 illustrates two potential deployments to aforementioned example devices (including AAS server components, deployment interactions, a broker server and stereotypes from the UMLsec/security profile).

The AAS of this component is represented by EcsAAS, actually the resources sub-model already mentioned in Section 3.7. This sub-model consists of the `ResourceUnit` instances (corresponding to single ECS runtimes) representing the resources and the installed/running `Container` instances. The `ResourceUnit` offers the operations to manage containers on the respective resource. Moreover, `ResourceUnit` is extended by service operations if the resource offers a `ServiceManager` (either installed in the same or in a different container on the same resource) as discussed in Section 3.7.

As for the generic IIP-Ecosphere components, regression tests validate the basic operations of the ECS runtime, i.e., an artificial test container manager and its AAS. For the Docker-based container management, the regression tests utilize a small Open Source container image and exercise the implemented operations. Akin to services, currently advanced container operations such as update and migration are not implemented.

Experiments with containers and AAS indicate that properties and operations work as expected. Simple operations can be executed in at maximum 5 ms runtime (or significantly less for monitoring properties as discussed in Section 3.7.3). Complex operations, e.g., starting a container depend on the time that is required by underlying operation of the container implementation, e.g., Docker. Direct execution on a non-virtualized operating system was not necessarily better in this regard. However, this experience strongly depends on the AAS and protocol implementation and, thus, is not representative.

We also validated starting Docker containers via the ECS runtime and the container manager, running the ECS runtime directly on the underlying operating system as well as running the ECS runtime in a Docker container. For the latter, a Docker-out-of-Docker (DooD)⁹¹ setup is required. Moreover, to achieve a certain genericity of the ECS runtime container, it is advisable to mount the containers via a folder of the host operating system into the ECS runtime container. The functions of the container management were validated on a Linux virtual machine running on a VMWare ESXi server as well as on the Phoenix Contact ACX 3152 mentioned in Section 3.7. As the container operations require a certain execution time, the minimum overhead created by an AAS-based management operation is not relevant here. The service capabilities are validated in several examples, most under regression testing, as well as some public demonstrators (cf. Section 6.6). The service/server functionality is currently subject to regression testing as well as the preparation of upcoming public demonstrators.

It is important to mention, that the sizes of the Docker container depend on the platform and application services that are installed. An ECS runtime with a DooD setup requires a container of around 1.1 GByte size (packed image of 444 MBytes), a service manager demands 509 MBytes (336 MBytes packed image) and a combined installation of ECS runtime and service manager into one container 600 MBytes (286 MBytes packed image). All containers can be installed and executed successfully even on an ACX 3152, typically with the platform server and the central broker installed on a server, e.g., the Linux virtual machine mentioned above. The running containers in idle mode allocate roughly 200 MBytes main memory (1.4 GBytes remain free on the ACX 3152), although at least 3 JVMs (ECS runtime, Service Manager and a local broker for the services) are running. If a simple service chain with 2 services is started, further 400 MBytes are allocated by one JVM per service, i.e., roughly 800 MBytes to 1 GByte memory remains free on the ACX 3152. Here, dependent on the actual load and service demands, we allow for some optimizations, e.g., to combine services with process backends, e.g., Python into the same JVM (ensemble services) or to limit the maximum memory allocation of the involved JVMs adequately. For the latter, the platform configuration model allows

⁹¹ <http://tdongsi.github.io/blog/2017/04/23/docker-out-of-docker/>

settings for the platform services as well as for individual application services, which are turned into executable artifacts by code generation (cf. Section 6).

Table 13: Review of realized⁸⁰ requirements for the ECS runtime component.

Requirement	Summary
R23	The ECS runtime provides the basis for dynamic deployment units. The actual deployment units are packaged during code generation for the configuration model. <i>Automatic creation of containers was started but the realization is not completed.</i>
R24	Generic and specific implementations of the ECS runtime aim at supporting the deployment to different types of resources/hardware. So far, we provide a generic ECS runtime in Java with a default container manager for Docker.
R25	Common functionalities of the resources (as well as service operations) are defined in the AAS of this component. The AAS also reflects the monitored resource meters of an ECS runtime instance.
R25a	The AAS of a resource is realized by the IIP-Ecosphere ECS runtime.
R25b	The AAS can and shall be deployed to a central platform AAS, in particular to integrate with the service management operations. <i>This may have to be complemented with a registration function in the resource/device management (cf. Section 3.8.2).</i>
R25e	Basic static properties are provided such as resources hosting a container.
R25f, R103a	The AAS also reflects the monitored resource meters of an ECS runtime instance. More specific meters can be added through the system metrics plugin.
R25g	Resource AAS (<code>ResourceUnit</code>) defines functions for the deployment.
R25h	<i>Resource AAS (<code>ResourceUnit</code>) does contain functions for exchanging deployment units at runtime, but the functionality is currently not implemented.</i>
R26	Deployment to on-premise resources is supported by the ECS runtime.
R27	<i>Optional deployment to connected IIP-Ecosphere instances is intended for the third development stream and, thus, currently not realized.</i>
R28	<i>Optional deployment to cloud resources, e.g., Google Cloud or Gaia-X is considered in the third development stream and, thus, this requirement is currently not realized.</i>
R29	The deployment unit provides an explicit interface in terms of an AAS.
R29a	Functional interfaces as well as quality properties of a deployment unit are provided via AAS.
R29b	The deployment unit AAS is linked to the resource sub-model and the services sub-model.
R29c	Contained services/containers are available through the deployment unit AAS.
R30	A deployment unit is encapsulated as container, in particular the services are encapsulated in artifacts to be deployed individually into containers.
R30a	Deployment units on IT level shall be technologically uniform, through the general use of Docker containers. The <code>ContainerManager</code> supports the exchange of the respective implementation/integration.
R30b	Deployment units on OT level can be technologically different, but we aim for Docker as the default technology. The <code>ContainerManager</code> supports the exchange of the respective implementation/integration.
R30c	The platform supports the external container repositories as full container name in the container descriptor. A container descriptor is specified via a URL, which may also point to an external descriptor location and further on to an external container repository.

Requirement	Summary
R31	A container shall contain only the required components/services as discussed in Section 3.7. <i>This depends on the packaging, for which an automated approach is not part of this release.</i>
R31b	Artifacts may contain optional components, which are then not executed as discussed in Section 3.7. <i>This depends on the packaging, for which an automated approach is not part of this release.</i>
R31c	Components/services in a container may be exchanged dynamically as supported by the service management in Section 3.7. <i>However, this functionality is not part of this release.</i>
R32	Container can contain data/models. These are artifacts in the containers and respective (typed) parameters are offered by the services cf. Section 3.7. <i>So far, the packaging of further resources is not realized.</i>
R33	A container can contain local data stores (in terms of services cf. Section 3.7). <i>Data stores are not part of this release.</i>
R35	AAS operation/access speed in containers was around 5 ms. <i>A sampling rate of 2 ms through a container is plausible, but must be shown in future experiments.</i>
R36	<i>Configuration of resources is part of device management, currently not completely realized and intended for a later development stream.</i>
R36a	<i>Configuration of resources is part of device management, currently not completely realized and intended for a later development stream.</i>
R36b	<i>Configuration of resources is part of device management, currently not completely realized and intended for a later development stream.</i>
R37	<i>Remote maintenance of resources is part of device management, currently not completely realized and intended for a later development stream. This can potentially be integrated with container orchestrator operations.</i>
R38-R44	<i>Security mechanisms are indicated in the architecture but not part of the platform implementation of this release.</i>
R45-R68	<i>Data protection mechanisms are indicated in the architecture but not part of the platform implementation of this release.</i>
R102	The platform allocates deployment units to available resources.
R105	Deployment happens via platform functions and AAS.

We conclude, that basic requirements for this layer are implemented. However, in comparison to the service management, connectors or transport component, in this component more advanced functionality is dependent on the automatic creation of containers or the device management. These components are scheduled for future releases.

3.8.2 Device/Resource Management

The device management shall support and ease the administration of devices, i.e., compute resources. As stated above, e.g. along with the ECS runtime in Section 3.8.1, the notion of devices in IIP-Ecosphere is rather broad as it involves edge, cloud and (on-premise) server devices. From a practical point of view, the scope includes all devices that potentially can run an ECS runtime (including the IT infrastructure from [41]) and/or a Service Manager. Also, different forms of installation for an ECS runtime as discussed in Section 3.8.1 are subject to the device management. It is important to recall that following [41], Industry 4.0 field devices such as machines are out of scope for the IIP-Ecosphere platform.

From [13] we know that the main requirements for the device management (as summarized in Table 14) refer in particular to the "Device Description Store", the "Device Configuration Tool" and the "ECS runtime" introduced in [41]. This includes the abstraction of vendor dependencies (R25.a),

on/offboarding (R25a) and device management (R25b). Common management functions which are neither listed in [13] nor [41], e.g., mechanisms for human interactions (acknowledgements), management techniques such as device templates or import functions for "asset data providers" [41] are desirable, but also well covered by existing platforms [35]. Thus, in [13, 41] it was intentionally left open, whether the IIP-Ecosphere just focuses on the essential capabilities mentioned in [13, 41] or provides also additional useful capabilities. Please note that quality requirements regarding data processing time limits, e.g., soft realtime, do not apply to management operations of the device management.

Besides this freedom, there are requirements that also prescribe the design of the device management. One important requirement is R7 which requires the use of AAS for the interfaces of all layers/components in the IIP-Ecosphere platform. On the one side, the device management must take the information in the platform AAS on available resources into account and use the operations provided there to manage resources, i.e., this component can require its own operations in the resources sub-model elements collections described in Section 3.8.1. On the other side, the device management shall provide relevant own additional operations (such as onboarding, selection of device templates) to upper layers such as the user interface of the platform. Where adequate, these operations shall be parameterized with the resource identifier from the resources sub-model (cf. Section 3.8.1). The functionality of the device management is influenced by given information (through AAS events⁹² and polling, R11), but may also directly influence the resource sub-model elements collection, e.g., adding/removing devices (potentially requiring subsequent operations, e.g., shutdown/migration of containers or services).

Moreover, the device management must take the virtual character of the IIP-Ecosphere platform into account (cf. Section 3.1). Therefore, it is mandatory that the device management is able to operate on multiple AAS of the structure described in this document rather than on "just" a singleton AAS of the IIP-Ecosphere platform. This allows taking other IIP-Ecosphere platform instances as well as underlying mapped-in platform instances into account. However, it is important to understand that access to these further AAS may be restricted, e.g., management operations are not allowed to be executed. This may be represented in terms of missing operations or AAS access limitations⁹³.

Primarily, for the device management Java 1.8 compatible libraries shall be used, although this constraint may be relaxed for this component as it will be utilized in the central IT installation. Regarding security (R38-R44) or data privacy (R45-R68), this may include the exchange/installation of encryption keys or certificates during onboarding.

Table 14: Specific requirements for device management (in addition to the general requirements in Table 2, Table 3)

Requirement	Summary
R25	Resources must be represented as AAS
R25a	Properties/functions of the ECS runtime must be represented as AAS
R25b	AAS of available resources must be announced to the platform
R25c	Platform must manage the available resources
R25d	Platform can offer procedures to facilitate the resource management

⁹² If possible, the component may rely on change events of the AAS implementation. However, in BaSyx events are currently in realization and, thus, not yet reflected in the AAS abstraction introduced in Section 3.5. Thus, the component design shall foresee event-based change notifications as well as (potentially less efficient) polling/scanning of the respective AAS structures.

⁹³ Currently, security and access restriction mechanisms are not (fully) in place in BaSyx and, thus, not reflected in the AAS abstraction introduced in Section 3.5.

Requirement	Summary
R27	Optional deployment to connected IIP-Ecosphere instances, i.e., device management must be designed transparently and equipped with appropriate access mechanisms
R28	Optional deployment to cloud resources, e.g., Google Cloud or Gaia-X
R36	Optional configuration of resources
R36a	Writing of resource configuration
R36b	Reading of resource configuration
R37	Optional remote maintenance of resources
R40	The platform must provide the usual security mechanisms like RBAC and TLS.
R41	The security mechanisms shall be integrated with common directory services.
R43	Safety mechanisms must describe their quality properties and their callable functions
R45	The data platform must provide fair and lawful processing of personal data
R47-R49	For processing personal data, data subject, legitimate purposes, adequacy, and storage time must be specified
R50	The platform must identify different categories of personal data
R51	The platform must ensure deletion, blocking, and authorization of personal data
R55	The platform must provide ways to object to the direct marketing of personal data
R57	The platform must offer possibilities to object to the decision support based on the automated processing of personal data
R136a	The platform shall provide an update functionality for the ECS runtimes.

For the device management, further existing components could be helpful, e.g., an IoT device management approach (to be frontended by AAS), a secure console (R37) or a storage for binary images (R36a, R36b, R136a). A discussion of potential components in the scope of the requirements for the resource management is provided by Pidun in [31]. There, for the IoT device management, two components are discussed, namely DeviceHive⁹⁴ and ThingsBoard⁹⁵. The specific capabilities of ThingsBoard, such as software-over-the-air, firmware-over-the-air, the broad range of supported protocols, an Angular user interface that could help realizing an optional user interface for the IIP-Ecosphere platform as well as a higher development activity in the recent time led to the suggestion of integrating ThingsBoard as one alternative technology into IIP-Ecosphere. For the binary storage, MinIO⁹⁶ and OpenStack Object Store Swift⁹⁷ were compared. Here the support for the de facto standard S3 made the difference and MinIO was suggested in [31].

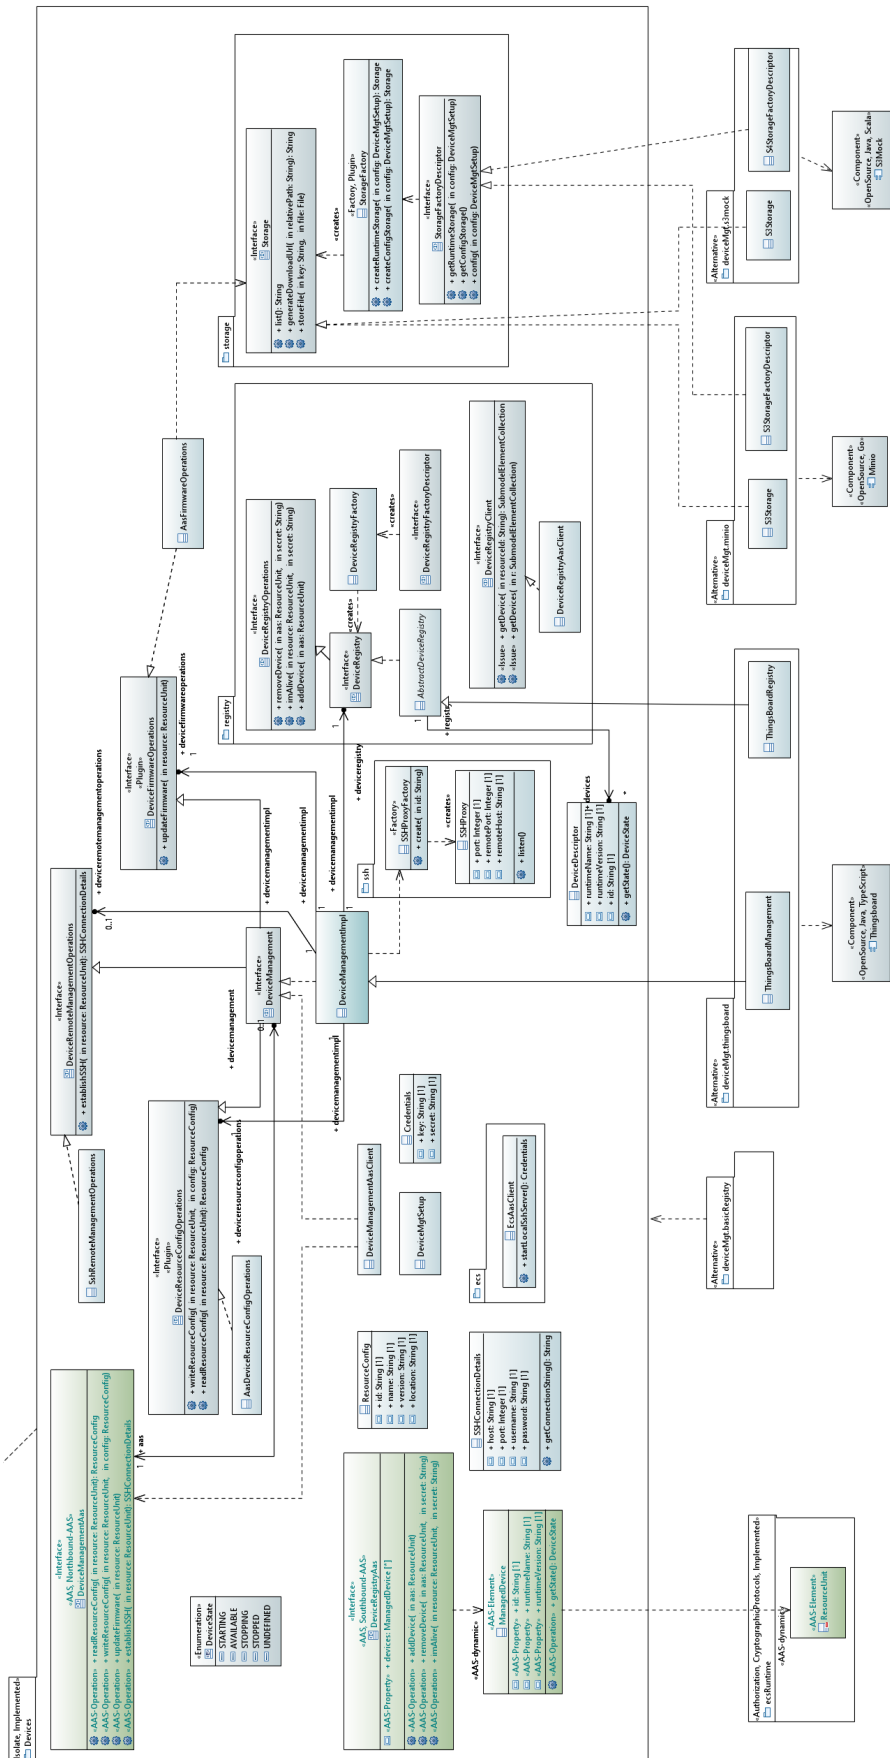
The architecture of this component follows the architectural suggestions in [31]. An overview is depicted in Figure 43. The component offers two AAS interfaces, a southbound interface in DeviceRegistryAas, and a northbound interface in DeviceManagementAas. The southbound interface is intended to enable a self-registration of devices and to notify the platform that they are available (heartbeat). This involves so-called ManagedDevice instances, which bridge between the ResourceUnit from Section 3.8.1 and specific information required by the underlying management approach, e.g., a different secondary device id. The northbound interface provides device information to higher-level components in the IIP-Ecosphere platform.

⁹⁴ <https://github.com/devicehive>

⁹⁵ <https://github.com/thingsboard/thingsboard>

⁹⁶ <https://github.com/minio/minio>

⁹⁷ <https://github.com/openstack/swift>



At the core is the DeviceManagement interface, which is composed of operation interfaces covering different aspects indicated by the requirements, such as resource configuration, remote management or firmware operations. The separation into different interfaces allows for a unified handling of implementation, AAS (client) and testing. The DeviceManagementImpl class unifies these interfaces (by delegation) and implements a default remote management approach via Secure Shell (SSH) based on temporary sessions created by request on trusted, registered devices. To rely on an existing, mature implementation, the communication is performed here via SSH streams rather than AAS (due to performance issues) or the transport layer (a pair of channels might be used for this purpose), although SSH may impose issues to Windows devices. Further parts of this component are the device registry abstraction and the storage abstraction. For both parts, the implementation is left open here and can be realized by alternative components, e.g., ThingsBoard, MinIO etc. As usual in IIP-Ecosphere, these parts define specific interfaces to the abstracted functionality, a JLS descriptor to create a concrete instance for the interface as well as supporting classes like AAS client implementations. A specific AAS client class (EcsAasClient) offers access to an extension of the ECS runtime from Section 3.8.1 to create a remote SSH endpoint on demand.

As indicated in Figure 43, for a specific setup, the device management offers the following alternatives:

- **ThingsBoard** as central management component. ThingsBoard ships with load balancing mechanisms, an own selection of internal protocol and frequently is setup via Docker containers. It must be installed separately in addition to server components.
- **basicRegistry** as a simple, in-memory implementation of the device registry interface. Can be used instead of ThingsBoard, but does not provide a user interface or persistency mechanisms, i.e., can be used for test setups.
- **MinIO** calls itself as the world's fastest object storage server. MinIO requires adequate setup on the server side. However, MinIO is licensed since April 2021 under AGPL, i.e., only an optional integration is permissible and it might not be the primary choice for IIP-Ecosphere installations.
- **S3Mock** is a Java/Scala-based object storage server for testing S3 implementations. In contrast to MinIO, it can be used without license limitations (MIT license), provides access via the Amazon S3 interface, but accepts any authentication, i.e., can be used for test setups. In contrast to MinIO, the S3Mock integration contains an IIP-Ecosphere lifecycle descriptor and, if the setup includes a storage server section, starts also a local storage installation (on the central IT side).

It shall be noted that due to usage of the well-known S3 protocol/interface for the object storage (although the individual technical interfaces of S3Mock and MinIO differ), the object storage integrations can act as storage connectors to storages located in a cloud and accesses can be directed to a cloud if stated in the component setup.

The device management component supports a simple on/offboarding process, *currently without manual approval of the operations*. If explicit on/offboarding is enabled (by default, this is currently not the case to ease development), a device must be explicitly on-boarded or off-boarded. This may lead to the exchange/removal of security certificates or encryption keys. On devices, that were not on-boarded, the platform may not execute operations. *Neither exchange of security information nor denial of operations are currently implemented.*

The functionality of the device management has been validated through many fine-grained test cases, see also [31]. There, the performance of the direct execution of individual device management operations using the ThingsBoard device registry and the MinIO S3 connector have been measured

and take in average 8-170 ms. If the operations are executed via the device management AAS sub-models, the operations take in average 11-204 ms.

Table 15 reviews the realized requirements for the device management components.

Table 15: Review of realized⁸⁰ requirements for the device management component.

Requirement	Summary
R25	Resources are represented as AAS (cf. Section 3.8.1). Managed devices are represented by their own AAS entities.
R25a	Properties/functions of the ECS runtime must be represented as AAS (cf. Section 3.8.1). An extension for device management is provided.
R25b	AAS of available resources must be announced to the platform through the device management onboarding process.
R25c	Platform manages the available resources.
R25d	<i>Platform can offer procedures to facilitate the resource management. Currently, no additional management mechanisms like templates are realized.</i>
R27	Optional deployment to connected IIP-Ecosphere instances. <i>Other platform instances can hook in although a platform instance classification is currently not provided.</i>
R28	Optional deployment to cloud resources can be realized through remotely installed ECS runtime and via the S3 storage connectors. <i>This requirement is currently not validated.</i>
R36	Resources can be configured through the device management.
R36a	Writing of resource configuration is supported by the device management.
R36b	Reading of resource configuration is supported by the device management.
R37	Optional remote maintenance of resources is supported by the device management through its SSH approach.
R40	The platform must provide the usual security mechanisms like RBAC and TLS. TLS depends on the installation of the AAS, <i>RBAC is currently not supported, only simple authentication.</i>
R41	The security mechanisms shall be integrated with common directory services, <i>which is currently not implemented.</i>
R43	<i>Safety mechanisms must describe their quality properties and their callable functions</i>
R45	<i>The data platform must provide fair and lawful processing of personal data</i>
R47-R49	<i>For processing personal data, data subject, legitimate purposes, adequacy, and storage time must be specified</i>
R50	<i>The platform must identify different categories of personal data</i>
R51	<i>The platform must ensure deletion, blocking, and authorization of personal data</i>
R55	<i>The platform must provide ways to object to the direct marketing of personal data</i>
R57	<i>The platform must offer possibilities to object to the decision support based on the automated processing of personal data</i>
R136a	The platform shall provide an update functionality for the ECS runtimes. This is <i>partially supported through the S3 store but currently not validated against the ECS runtimes.</i>

3.8.3 Monitoring

Service execution shall be monitored, in terms of resources but also in terms of functionality, e.g., through application specific probes and alerts. Therefore, the IIP-Ecosphere platform foresees a set of generic built-in monitoring probes (cf. Section 3.7) as well as application-specific probe extensions that communicate their information via topic streams to one or multiple monitoring information aggregators. In turn, aggregators provide their state to upper level layers. Also (application-specific) alarming via specific streams shall be supported. In addition to the service monitoring, the IIP-

Ecosphere platform shall also monitor resources via the installed ECS runtimes and also the execution of the ECS runtime.

While the probing of the individual services or ECS runtimes/resources happens on the devices (and thus belongs to Section 3.7 or Section 3.8.1, respectively), the main task of this component is to aggregate the information on IT infrastructure level (see also [41]). The aggregation of the received values shall follow existing guidelines, approaches, relevant standards or norms in Industry 4.0. As the IIP-Ecosphere platform shall operate across a plethora of resources (and connected or underlying platforms and their resources, if available), the monitoring component shall foresee (optional) hierarchical aggregation to distribute the input load and to increase the efficiency.

Table 16 summarizes the basic requirements for the monitoring component from [13]. These focus on devices/resources, services and alarming/alerts, in generic or application-specific fashion (e.g., through specific monitoring services hooked into the data processing chain). Akin for the device management, one important general requirement is R7 which requests the use of AAS for the interfaces of all layers/components in the IIP-Ecosphere platform. On the one side, the monitoring must take the information in the platform AAS on available resources into account and use the information provided by services and resources through their local monitoring. This component may require further properties than those, e.g., described in Section 3.8.1. On the other side, the device management shall provide relevant (aggregated) information and own operations to upper layers such as the user interface of the platform. The functionality of the monitoring component shall rely on underlying information (through AAS events⁹² and polling, R11) in the services/resource sub-model elements collections. However, due to potential performance issues of the AAS approach, for urgent alarms/alerts also a second path via the Transport component is more adequate.

As described for the device management, the monitoring component must take the virtual character of the IIP-Ecosphere platform into account (cf. Section 3.1). Therefore, it is mandatory that the monitoring is able to operate on multiple AAS of the structure described in this document rather than on “just” the singleton AAS of the IIP-Ecosphere platform. This allows taking other IIP-Ecosphere platform instances as well as underlying mapped-in platform instances into account. However, it is important to understand that access to these further AAS may be restricted, e.g., access to information is limited. This may be represented in terms of missing properties or AAS access limitations⁹³.

Table 16: Specific requirements for monitoring (in addition to the general requirements in Table 2, Table 3)

Requirement	Summary
R4d	Execution of services must be supervised/monitored.
R4e	Service monitoring shall be parameterizable.
R4f	Service monitoring shall be realized by application-specific services
R25e	The AAS of the resource must describe static properties of the resource
R25f	The AAS of the resource must describe dynamic properties of the resource
R46	The collection of personal data must be for specified, clear and legitimate purposes
R53	The platform must provide a mechanism for notifications regarding rectification, deletion, blocking, and leakage.
R63	The platform shall provide a mechanism to capture user privacy and security requirements
R118a	The platform shall provide warning and alerting services to the AI toolkit.
R122	Based on the configuration model, the platform shall be able to decide on the optimal use of AI methods, i.e., reflection of monitoring information into the configuration model is needed.

Primarily, for the monitoring component Java 1.8 compatible libraries shall be used, this constraint can be relaxed for the central monitoring components. As first (alternative) monitoring component we decided for an integration of the Prometheus service and resource monitoring approach (open source, Apache License). Prometheus is based on gathering data from HTTP/REST servers exposing monitoring endpoints, allows for configuring evaluation rules on the gathered data, stores the data in a time series data base and exposes the aggregated information again as HTTP/REST service endpoints (including an alert manager). However, direct HTTP access across all resources in a production system may not be permitted, i.e., some intermediary representation might be required.

One approach could rely on directly reading out the platform AAS, the devices and services AAS or submodels. Initial experiments using BaSyx as backend for a web-based UI (cf. Section 3.14) are promising. Here, standardized submodels (as started by the IDTA for a resources submodel) could lead in the future to existing, reusable components for AAS based scraping of monitoring information, and, thus, could ease the effort of integrating such monitoring systems.

An alternative approach, as illustrated in Figure 44, is to provide the monitoring information through the transport layer as envisioned in [4]. For Prometheus, this is then similar to approaches allowing Prometheus to monitor resources over network borders, where a firewall or a gateway provides a proxying service⁹⁸, which collects all relevant information in the subnet on behalf of Prometheus and offers the collected information on individual endpoints to the scraping process of Prometheus. Although this requires an additional server process in the Prometheus case, it also allows for the flexible integration of other monitoring systems, as the data is provided independently and just must be translated into a format that can be understood by the respective monitoring system, e.g., a transport-to-MQTT translation that feeds information into the monitoring component.

In our case, the integration of Prometheus into the IIP-Ecosphere platform receives the monitoring data of individual resources via transport communication, exposes this information in an own (local) web server and adjusts the Prometheus configuration so that new devices are considered for scraping via the (local) web server. The implementation of such a metrics exporter is prepared in the generic IIP-Ecosphere monitoring component, while the Prometheus-specific integration is done in the alternative Prometheus integration component. The Prometheus integration component also contains the Prometheus binaries for Windows and Linux as well as an own lifecycle descriptor that starts and shuts down Prometheus. Within this lifecycle, a bridging metrics exporter as well as an alert monitor are started. We use the Prometheus alertmonitor⁹⁹ (Apache License) as the Prometheus client library does not provide support for alerts. The alertmonitor scrapes the alert manager HTTP API of Prometheus in regular fashion and turns alerts into alert instances of the transport layer (alert stream, see Table 25 in Section 8.1). The setup of Prometheus is defined in the configuration model and the setup information is generated during the platform instantiation process. Part of this setup is also the information, whether we rely on an installed Prometheus server or whether we have to start the included binaries.

The monitoring component defines an own AAS submodel, which currently consists of a list of recent alerts. Individual, aggregated monitoring values as well as changes of monitoring rules will be subject of the next releases.

Ultimately, the Prometheus monitoring component is configured and integrated through the configuration model and the platform instantiation. However, realizing the bridging approach requires careful handling of the embedded Tomcat webserver, as otherwise in particular our default test broker Apache Qpid may throw `NullPointerExceptions` when the Tomcat instance creates server

⁹⁸ <https://github.com/pambrose/prometheus-proxy>

⁹⁹ <https://github.com/matjaz99/alertmonitor>

contexts. This would significantly limit the flexibility of choosing transport protocols in the platform. The current approach works with a single root context, into which a default metrics servlet and one servlet per device in an IIP-Ecosphere platform installation are added. The default metrics servlet is added statically when the server instance is created, the device servlets upon status messages of the platform. Further evaluation also of the performance will be subject to further work / releases.

Some components conflict regarding their dependency versions, e.g., in the past BaSyx and Spring regarding the required versions of the Servlet container Tomcat, partially running both in the same JVM was not possible. For this purpose, the platform instantiation creates an own monitoring component that can be executed besides the platform services in an own process. Moreover, the Prometheus monitoring component defines three lifecycle profiles, the complete Prometheus integration, only Prometheus or only the integration. In one of the next releases, we may abandon the creation of an own component and rely only on the lifecycle profiles.

The monitoring of system-level meters and application level meters (data items received/sent) has been validated through the Prometheus UI. Added resources (ECS-Runtime/Service Manager) when they occur are taken up, system- and application-level meters are categorized according to device id (and service id for application-level) and displayed individually. Aggregated values or rates can be calculated from this information on Prometheus level. Currently, the underlying approach based on micrometer automatically adds several technical system-level meters. Moreover, Spring also adds additional meters. Most of these meters are not relevant on platform level and could be filtered out. Also per monitoring diagram, currently all meter information is transmitted, while updates could focus on changing values and omit already known static values like descriptions or monitoring units. Both improvements are out of scope for this release, but planned for the next release of the IIP-Ecosphere platform.

Table 20 summarizes the state of realizing the requirements for the monitoring approach.

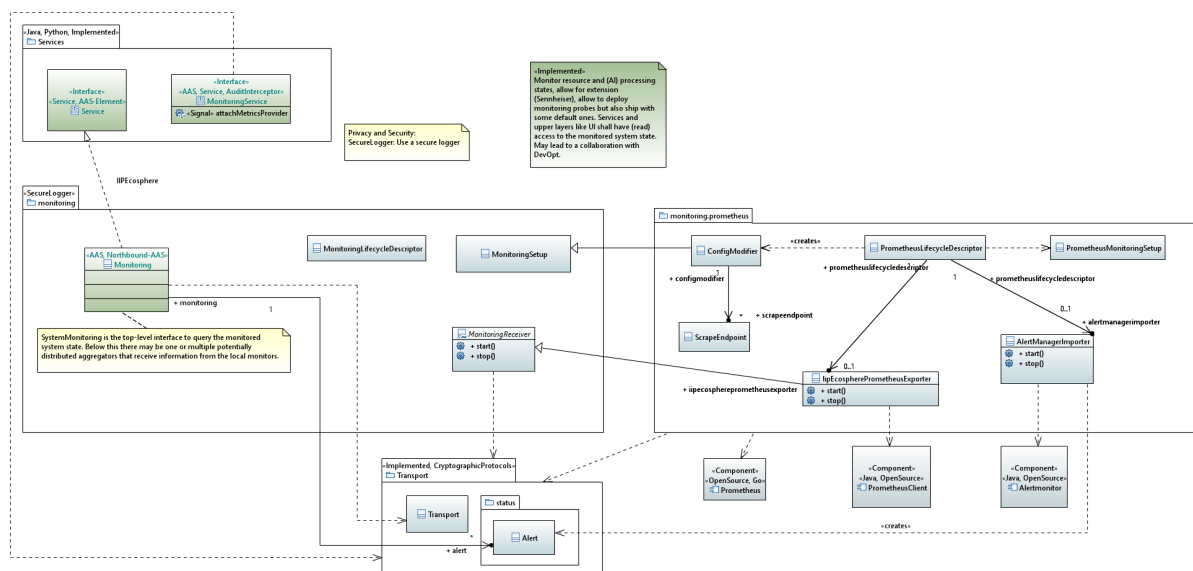


Figure 44: Monitoring (comments cropped)

Table 17: Review of realized⁸⁰ requirements for the monitoring component.

Requirement	Summary
R4d	Execution of services is be supervised/monitored.
R4e	Service monitoring is parameterizable via monitoring rules and application-specific monitoring services, which can offer own parameters.
R4f	Service monitoring shall be realized by application-specific services. Generic monitoring probes, e.g., for throughput are integrated automatically by the platform instantiation. Further generic resource monitoring probes are provided by used libraries, e.g., Spring or the Prometheus Java client. The interface of probe services is defined and implementations can be based on the alert format defined by the transport component.
R25e	The monitoring component provides an own AAS, which builds up a dynamic structure. Currently, the focus of the AAS is on monitoring alerts <i>rather than aggregated monitoring values</i> .
R25f	The monitoring component provides an own AAS, which builds up a dynamic structure. Currently, the focus of the AAS is on monitoring alerts <i>rather than aggregated monitoring values</i> .
R46	The monitoring approach does not collect personal data.
R53	The monitoring approach does not collect personal data.
R63	<i>The platform shall provide a mechanism to capture user privacy and security requirements.</i>
R118a	The interface of probe services is defined and implementations can be based on the alert format defined by the transport component.
R122	<i>Currently, the monitoring information is not reflected into the configuration model.</i>

3.9 Storage, Security and Data Protection Layer

The Storage, Security and Data Protection Layer is responsible for managing security aspects of the IIP-Ecosphere platform based on the platform configuration, for offering security-enhancing services (such as anonymization or pseudonymization) as indicated in Figure 45, but also for secure integration of (encrypted) data lakes or clouds. As discussed in Section 3.1, the purpose of this layer is not to realize typical cross-cutting security mechanisms, which will be subject to the security discussion in Section 7.

Table 22. summarizes the specific requirements for the Storage, Security and Data Protection layer. We do not focus on the configuration aspects (R40a, R40b, R41a, R42, R44, R64a, R65a) here, as we do so later in the discussion of the Configuration Layer in Section 3.11.

Table 18: Specific requirements for the Storage, Security and Data Protection layer
(in addition to the general requirements in Table 2, Table 3)

Requirement	Summary
R38	Use of services only for authorized persons.
R39	Personal data is only changed by authorized persons.
R40	Common security mechanisms, such as RBAC or TLS.
R41	Integration with common directory services, e.g., LDAP
R41b	If no directory service is available, a mechanism for managing user accounts must be provided.
R43	Safety mechanisms must describe their (quality) properties and callable functions as AAS.
R43a	The AAS shall describe the respective impact on the performance.
R43b	Selected security mechanisms shall be deployable.
R45	The platform must provide fair and lawful processing of personal data.
R46	The collection of personal data must be for specified, clear and legitimate purposes.

Requirement	Summary
R47	The platform shall avoid the processing of personal data as much as possible.
R47a	Applications running on the platform shall avoid processing personal data as much as possible.
R47b	Models used on the platform shall be protected against privacy attacks.
R48	Store personal data in a form that permits identification of data subjects for no longer than is necessary for the purposes for which the data is processed.
R49	Process personal data adequate and relevant to the purpose and limited to what is necessary for the purposes of the processing.
R49a	Applications shall process personal data adequate and relevant to the purpose and limited to what is necessary for the purposes of the processing.
R50	Identify different categories of personal data processing.
R51	Ensure the authorization, deletion or blocking of personal data.
R52	Store personal data in a structured, common and machine-readable format.
R53	Notification mechanism for rectification, deletion, blocking, leakage of personal data.
R54	Ensure possibilities to object to the processing of personal data.
R55	Provide ways to object to the direct marketing of personal data.
R56	Provide options to object to the transfer of personal data to third parties.
R57	Offer possibilities to object to decision support based on the automated processing of personal data.
R58	Support the detection of personal data breaches and their reporting to data subjects.
R59	Ensure authorization with regard to access to personal data.
R60	Provide privacy principles when consent is required for the processing of personal data.
R61	Enable users to control their personal data requirements.
R62	Facilitate the assessment of data protection impact assessment to identify threats and risks in the processing of personal data.
R63	Provide a mechanism to capture user privacy and security requirements.
R64	The privacy mechanisms shall anonymize the information in specified data fields.
R65	The privacy mechanisms shall detect and anonymize personal data contained in free text.
R66	Pseudonymization shall keep pseudonyms in the system only for the time actually needed.
R66a	After the time of need has expired, a new pseudonym shall be assigned to the same person.
R66b	The purpose and scope shall be aligned with the operating agreement when used internally.
R66c	Pseudonymization shall be applied in case of external use only after the consent of the person concerned.
R67	The platform shall capture and classify generated cookies or similar identifiers stored on end devices.
R68	The platform shall provide the possibility of automatic deletion of such (R61) identifiers, as well as deletion at the request of the user.

3.9.1 KODEX platform service

The privacy enhancing service in this layer integrates the KODEX privacy and security engineering toolkit¹⁰⁰ by KIPROTECT into the IIP-Ecosphere platform. Currently, we focus on two alternative

¹⁰⁰ <https://heykodex.com/>, <https://github.com/kiprotect/kodex>

integration forms, one via command line streams, i.e., as shown in Figure 45 (AbstractStringProcessService from the Java service environment) as well as a REST-based alternative. It is important to mention that KODEX is a generic tool that requires some form of setup to operate on the incoming data in the intended manner. Thus, KODEX acts here as a blueprint for rather generic data services in the IIP-Ecosphere platform. It is interesting to mention that KODEX is realized in GO, i.e., not in Java, and it was the first external data processing service integrated into the IIP-Ecosphere platform. To cope with the genericity of KODEX, some design decisions were made for KODEX that apply analogously to other external services:

- The KodexService is parameterized over the incoming and outgoing data types. To transfer data instances correctly to KODEX, respective type translators (the more generic form of a Serializer) are required. These type translators shall be provided by the utilizing code, e.g., a Spring Cloud Service node generated from the configuration model, which collects the knowledge about incoming and outgoing types of all service chains of all applications on a certain IIP-Ecosphere instance.
- The customization of KODEX happens in terms of certain files that specify the data model. Akin to the type translators, the contents of these files are determined upon integration into a service processing chain and shall be generated from the configuration model. These files shall be packaged into a ZIP archive (named according to the using node in the service chain) and stored in the service implementation artifact as specified in the respective process part of the service deployment descriptor, which is also contained in the service implementation artifact. When starting the service node, the deployment descriptor is consulted, the artifacts are extracted and the customization files are placed into the home directory of the process implementing the service, here KODEX.
- In the (extracted) home directory of the process, also the service implementation must be located, i.e., in the KODEX case the operating-system specific binaries. Such implementation files shall be packaged into a “binary” Maven ZIP artifact and deployed along with the service integration code, here KodexService. When integrating the generic service code into a service chain, the Maven identification of the service implementation, here KODEX, is known, and so is the deployed implementation (here binaries for different operating systems). During automated instantiation/integration, the “binary” ZIP is packaged into the service implementation artifact and named respectively so that it can be extracted upon service start along with the customization files as described above.
- Upon code generation of the Spring nodes, further customizations may happen, e.g., service-specific customization files could be created.

Further integrations, e.g., using the REST API of KODEX as well as a performance comparison among different forms of the integration are currently being conducted. Initial results confirm that REST outperforms command line streams on Windows. For example, processing a batch of 1000 tuples, windows with command line streams takes 15 ms per tuple in average, REST on Windows 0.22 ms, command line streams on Linux 1.4 ms and REST 2 ms on Linux.

Regarding licenses, it is important to mention that at the point in time of writing this document, KODEX is licensed under AGPL. However, viral AGPL rules do apply to binary code, i.e., using the KODEX binaries with respective credits does not taint the license limitations of IIP-Ecosphere. Moreover, KODEX is only integrated, if it is explicitly used as a service in an application and only becomes active when the respective service chain is started.

3.9.2 Data lakes / Data bases

Work on data lakes has just been started, i.e., the KODEX integration is currently the only component that is contributed by the Storage, Security and Data Protection Layer to this release. Precursor work

in BMBF HAISEM-Lab on the performance of various time series data bases (e.g., Influx, Timescale, etc.) as well as discussions with partners on their experiences gives an indication for potential candidates to be integrated.

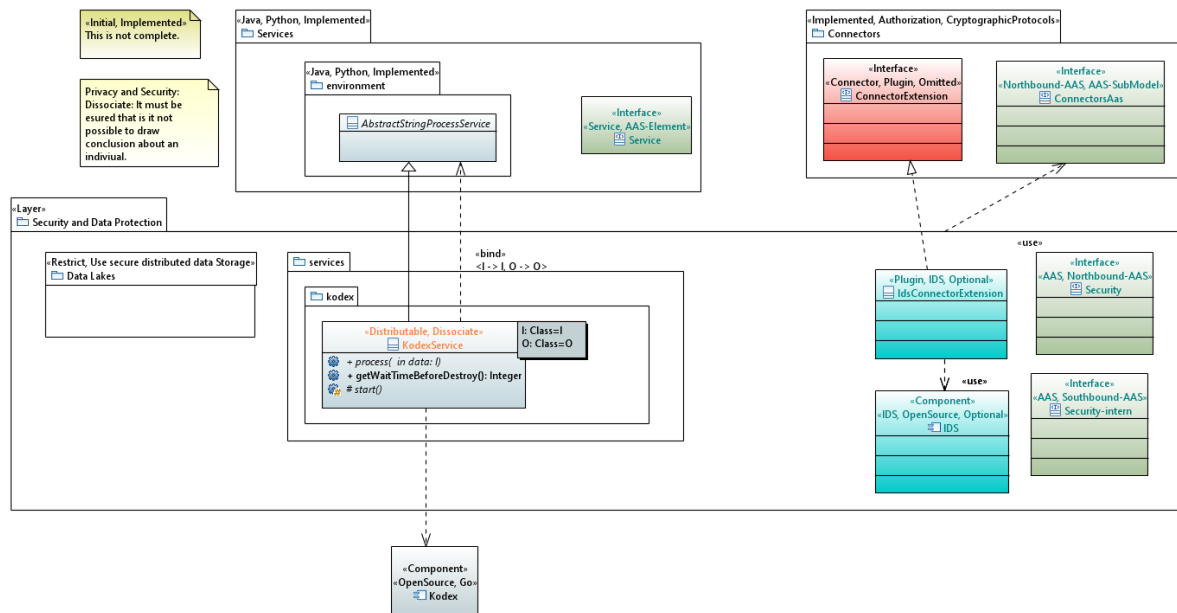


Figure 45: Overview of the data and security (comments cropped)

Table 19 summarizes the realized and open requirements for the Storage, Security and Data Protection Layer. As indicated in that table, many of the security and privacy requirements are still open while some are fulfilled, e.g., through the integration of KIPROECT KODEX.

Table 19: Review of realized⁸⁰ requirements for the Storage, Security and Data Protection Layer.

Requirement	Summary
R38	<i>So far, the platform does not provide authorization mechanisms.</i>
R39	<i>So far, the platform does not provide authorization mechanisms.</i>
R40	<i>Some components like the transport connectors exhibit settings for enabling TLS. This is currently not managed through the Storage, Security and Data Protection Layer, the configuration or the device management. Currently no mechanism for RBAC exists.</i>
R41	<i>Currently, there is no integration with common directory services, e.g., LDAP.</i>
R41b	<i>Currently, there is no mechanism for managing user accounts.</i>
R43	<i>Currently, there is no AAS for this layer.</i>
R43a	<i>Currently, there is no AAS for this layer.</i>
R43b	<i>Selected security mechanisms shall be deployable, e.g., the KODEX service.</i>
R45	<i>Currently there are no mechanisms for provide fair and lawful processing of personal data.</i>
R46	<i>Currently, there are no mechanisms to handle that the collection of personal data is only for specified, clear and legitimate purposes.</i>
R47	<i>The platform shall avoid the processing of personal data as much as possible.</i>
R47a	<i>The design of applications is out of the scope of the platform.</i>
R47b	<i>Models are currently not protected against privacy attacks.</i>
R48	<i>Currently, data is not stored.</i>
R49	<i>Currently, there are no mechanisms to ensure that personal data is processed adequate and relevant to the purpose and limited to what is necessary for the purposes of the processing.</i>
R49a	<i>The design of applications is out of the scope of the platform.</i>

Requirement	Summary
R50	<i>Currently, no categories of personal data processing are identified.</i>
R51	<i>Currently, no authorization, deletion or blocking of personal data is supported.</i>
R52	<i>Currently, no data is stored.</i>
R53	<i>Currently, there is no notification mechanism regarding rectification, deletion, blocking, leakage.</i>
R54	<i>Currently, there are no possibilities to object to the processing of personal data.</i>
R55	<i>Currently, there are no ways to object to the direct marketing of personal data.</i>
R56	<i>Currently, there are no options for objecting to the transfer of personal data to third parties.</i>
R57	<i>Currently, there are no possibilities to object to the decision support based on the automated processing of personal data.</i>
R58	<i>Currently, there is no detection of personal data breaches and their communication with data subjects.</i>
R59	<i>Currently, there is no authorization mechanism.</i>
R60	<i>Currently, there are no supported privacy principles when consent is required for the processing of personal data.</i>
R61	<i>Currently, there are no means to enable users to control their personal data requirements.</i>
R62	<i>Currently there is no assessment of data protection impact assessment to identify threats and risks in the processing of personal data.</i>
R63	The KODEX service is one mechanism to capture user privacy and security requirements.
R64	The KODEX service anonymizes information in specified data fields.
R65	The KODEX service can anonymize personal data contained in free text.
R66	<i>Pseudonymization shall keep pseudonyms in the system only for the time actually needed.</i>
R66a	<i>After the time has expired, a new pseudonym shall be assigned to the same person.</i>
R66b	<i>The purpose and scope shall be aligned with the operating agreement when used internally.</i>
R66c	<i>Pseudonymization shall be applied in case of external use only after the consent of the person concerned.</i>
R67	<i>The platform currently does not store cookies or similar identifiers stored on end devices. This may apply to an extension of the device management.</i>
R68	<i>Currently, the platform does not provide the possibility of automatic deletion of such (R61) identifiers, as well as deletion at the request of the user.</i>

3.10 Reusable Intelligent Services Layer

On top of the layers discussed before, the Reusable Intelligent Services Layer provides AI mechanisms in reusable and configurable manner and integrates received/monitored data with additional information such as product order information or floor plans to provide further valuable input to the AI. In this section, we briefly discuss the specific requirements (Section 3.10.1), the integration of RapidMiner RTSA as AI platform service (Section 0), further service candidates ahead (Section 3.10.4) as well as the planning for an AI toolkit (Section 3.10.6).

3.10.1 Specific Requirements

This section lists the specific requirements for the Reusable Intelligent Services Layer. It is important to remember from [13] that the actual realization of this layer / the AI toolkit is not part of the responsibilities of the core platform team rather than of all IIP-Ecosphere partners, e.g., the partners in the AI-Accelerator. However, the work of the Think Tank Platforms provides the technical framing for the AI toolkit, which shall follow the requirements in Table 22.

*Table 20: Specific requirements for the Reusable Intelligent Services layer
(in addition to the general requirements in Table 2, Table 3)*

Requirement	Summary
R110	Definition of interfaces for relevant AI components in industrial production.
R110a	Interfaces must be realized as AAS.
R110b	AAS interfaces must describe functional and quality aspects.
R111	AI toolbox must be extensible
R112	Elements of toolbox shall be distributable.
R112a	Parameters of the elements shall be defined in the configuration model.
R112b	Properties of distribution shall be defined in the configuration model.
R112c	Distribution shall be subject to restrictions for individual elements.
R113	AI toolkit shall support AI components in common programming languages/environments.
R113a	Support for Python, Pandas/Numpy, Knime, Scikit-learn, Tensorflow and RapidMiner.
R113b	Support for a middleware for agent management based on the Industry 4.0 language.
R114	Provide relevant AI methods common in/suitable for production.
R114a	AI methods shall be generic, parameterizable and adaptable.
R114b	AI toolkit shall support Transfer Learning.
R114c	AI toolkit shall support Reinforcement Learning.
R114d	AI toolkit shall support simple statistical procedures for decision making.
R114e	AI toolkit shall support neural networks.
R114f	AI toolkit shall support time series classification.
R114g	AI toolkit shall support anomaly detection.
R114h	AI toolkit shall provide methods for state detection.
R115	AI toolkit shall support AI models in standardized formats.
R116	AI toolkit shall support pre- and post-processing of data.
R116a	Pre- and post-processing procedures shall be generic, parameterizable and customizable.
R116b	AI toolbox can provide feature design / digital signal processing services (bandpass filters, FTF transforms, running RMS)
R117	AI toolbox and the AI methods included shall provide transparency and traceability about their decision-making
R118	The platform must provide central services to the AI construction kit.
R118a	The platform shall provide warning and alerting services to the AI toolkit.
R118b	The platform must provide storage services to the AI toolkit.
R118c	The platform must provide security services to the AI toolkit.
R118d	The platform must provide privacy services to the AI toolkit. [
R118e	The platform must provide data integration services to the AI toolkit.
R119	Training of AI methods shall happen automatically and in parallel/background. The trained models shall then be offered to the components of the platform.
R119a	The release of the trained model shall be done manually by a user with appropriate rights.
R119b	The release of the trained model shall be determined via settings in the configuration model.
R119c	The release of the trained model can be automatic (if specified in the configuration model).
R119d	The changes triggered by a release of a trained model shall be traceable.
R119e	Changes initiated by a release shall be reversible.

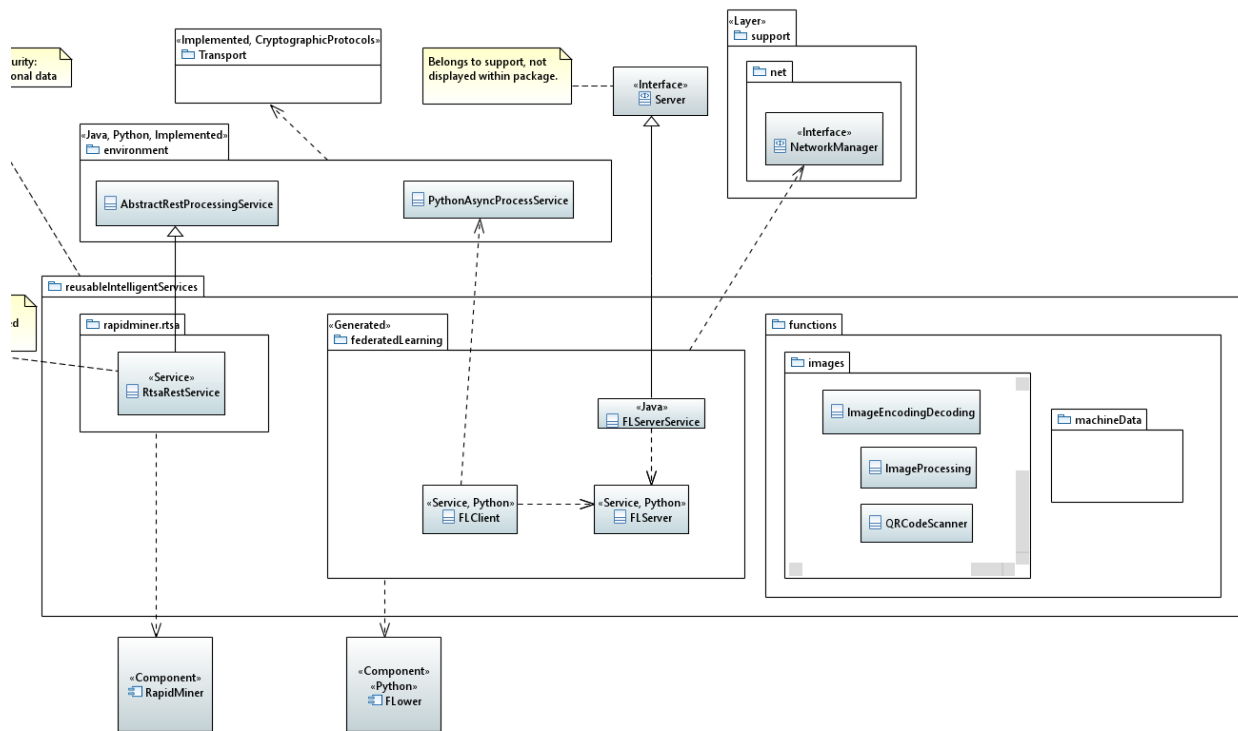


Figure 46: Reusable Intelligent Services and data processing function library (cropped)

3.10.2 Data Processing Function Library

Application of AI methods encompasses more than just the AI methods. Usually, also functionalities for data pre-processing etc. are required. As a basis to meet the requirements regarding pre-processing and data transformation functions (in [41], e.g., frequency analysis), we equipped the platform with a library of functions for data processing. In this version of the platform, first functions were integrated, namely

- Image transcoding from/to base64 strings.
- Image processing such as grayscaling, rescaling or thresholding.
- Barcode/QR-code detection based on the Java library *zxing*¹⁰¹ and, as optional fallback, the Python library *pyzbar*¹⁰². For the Python fallback, respective packages must be installed.

3.10.3 RapidMiner RTSA service

RapidMiner is a pioneering company in the fields of data analysis and AI. Their work on the RapidMiner Platform (in Java) and the RapidMiner Studio shows that AI composable from building blocks is not only a vision. One fundamental component in the RapidMiner ecosystem is the Real Time Scoring Agent (RTSA), a REST-based execution environment for deployments created by RapidMiner Studio.

Following the ideas in [41], a separation of data science exploration and design processes from the actual execution/deployment is desirable. Thus, RapidMiner is an excellent example for such an approach integrated into the IIP-Ecosphere platform. While the DataAnalyst can first create a data science process for given data in his/her own environment, the created process (as “deployment”) can later be deployed by the IIP-Ecosphere platform and executed under the control of the Service Manager on top of RTSA. For the integration, mainly the data input/output formats must match, i.e., the data provided by the IIP-Ecosphere platform (output of a connector/service) becomes the input

¹⁰¹ <https://zxing.org/w/decode.aspx>

¹⁰² <https://pypi.org/project/pyzbar/>

for the RTSA deployment and the output of the RTSA deployment becomes an input for upstream IIP-Ecosphere services.

Along these lines, the platform-supplied `RtsaRestService` (see Figure) integrates the RTSA and links the platform data streams to the input/output of the RTSA. The realization is similar to the REST-based integration of KODEX discussed in Section 3.9.1. In the configuration, a deployment file for the RTSA is specified, which is packaged by the application/platform instantiation along with the RTSA binary into the respective service artifact. Here it must be considered that RTSA is a commercial service, i.e., it requires a license file and cannot be distributed openly. For testing the service integration and the platform instantiation the test part of the component ships with an RTSA mockup (`FakeRtsa`), which acts as a REST server pretending to be an RTSA instance with a deployment. This fake RTSA can be configured in limited form to transform the input data, e.g., by changing fields or adding fields having a random number value. While this is not needed for the plain RTSA testing, it can help testing data flows for an application if a real RTSA is not available. For the platform instantiation, either the real or the mocking RTSA and its deployment are given in a specific folder according to the platform naming for binary files from which the instantiation process can take up the binaries.

3.10.4 Flower-based Federated Learning

One AI method that is of particular interest for the industrial production and for the integration within the IIP-Ecosphere platform is federated learning. Federated learning services act as services as they consume data and produce predictions, but they also share information about their AI model with other federated learning services (usually of the same application) so that the other services can learn new knowledge faster, e.g., acceptable anomalies. So far, also due to the client-server approach, federated learning is often viewed as a complex approach, which might be used more frequently in the industrial practice if an easy-to-use form of federated learning would be available. We aim at demonstrating one approach to easy-to-use federated learning in terms of the upcoming Hannover Messe 2023 IIP-Ecosphere demonstrator. Currently, the integration is still in development, but technically it is based on the service/server communication introduced in Section 3.7.4.

Figure illustrates the planned integration for the Flower¹⁰³ framework written in Python. Services (Federated learning client) and the server (usually assigned to one or multiple applications) are supposed to be specified in the configuration model. Both, client services and server service are realized in Python based on Flower and executed through the Python Service Environment and the respective Java integration counterpart (`PythonAsyncProcessService`). As explained in Section 3.7.4, in contrast to the client services, the server does not consume/produce regular service data streams. To be manageable within the platform, the server needs a Java counterpart which is also a `Server`, which manages the (hidden) service application lifecycle and relies in most of its functionality directly on `PythonAsyncProcessService`. All related parts, client service, server service and Java counterpart are supposed to be generated from the configuration model. While we are convinced that the server can be generated completely, we will probably generate a template for the client service (including the basic technology code for, e.g., tensorflow and numpy).

3.10.5 Service candidates ahead

Further service candidates are envisioned for integration, e.g., the demonstration of Python-based AI from the Hannover Messe 2022 demonstrator or individual service components from the IIP-Ecosphere Sennheiser demonstrator.

¹⁰³ <https://flower.dev/>

3.10.6 IIP-Ecosphere AI Software Service Concept

Currently, the partners are working on an AI service concept as foundation for the AI service toolkit. It will probably cover AI software services on different levels of granularity, ranging from atomic binding blocks over application sub-meshes up to entire application use cases. Some of these ranges can already be covered by the RTSA integration, individual points within this range of AI services will be filled as examples by the expected services ahead. Thus, we expect experience from the initial practical integration of AI services that can help to complement and complete the concept of the AI toolkit.

3.10.7 Requirements Discussion

We review now the realized requirements in Table 23.

Table 21: Review of realized⁸⁰ requirements for the Storage, Security and Data Protection Layer.

Requirement	Summary
R110	<i>Definition of interfaces for relevant AI components in industrial production depends on the AI service concept (Section 3.10.6).</i>
R110a	Service interfaces are managed and represented by the platform as AAS.
R110b	Service interfaces must express quality aspects as AAS.
R111	<i>AI toolbox must be extensible, which depends on the AI service concept (Section 3.10.6).</i>
R112	Elements will be realized as platform services and therefore may be distributable (depending on AI requirements).
R112a	Services in the platform can have parameters that are defined in the configuration model.
R112b	Properties of distribution shall be defined in the configuration model. <i>Currently, the configuration contains only the information whether a service is distributable.</i>
R112c	<i>Currently no distribution applies.</i>
R113	Supported through the service environments for Java and Python and the process-based service integration.
R113a	Support for Python and RapidMiner is provided. Pandas/Numpy, Knime, Scikit-learn, Tensorflow come through dependencies that can be specified in the configuration model. <i>The automated creation of containers with these dependencies is ongoing.</i>
R113b	<i>This depends on the AI service concept (Section 3.10.6).</i>
R114	<i>This depends on the AI service concept (Section 3.10.6).</i>
R114a	<i>This depends on the AI service concept (Section 3.10.6).</i>
R114b	<i>Currently a federated learning approach is in integration. Further approaches depend on the AI service concept (Section 3.10.6).</i>
R114c	<i>This depends on the AI service concept (Section 3.10.6).</i>
R114d	<i>This depends on the AI service concept (Section 3.10.6).</i>
R114e	<i>This depends on the AI service concept (Section 3.10.6).</i>
R114f	<i>This depends on the AI service concept (Section 3.10.6).</i>
R114g	<i>Currently a federated learning approach is in integration. Further approaches depend on the AI service concept (Section 3.10.6).</i>
R114h	<i>This depends on the AI service concept (Section 3.10.6).</i>
R115	<i>This depends on the AI service concept (Section 3.10.6).</i>
R116	<i>This depends on the AI service concept (Section 3.10.6).</i>
R116a	<i>This depends on the AI service concept (Section 3.10.6).</i>
R116b	<i>This depends on the AI service concept (Section 3.10.6).</i>
R117	<i>This depends on the AI service concept (Section 3.10.6).</i>
R118	<i>Already available central services are provided.</i>
R118a	The transport layer allows for warnings and alerts.

Requirement	Summary
R118b	So far, no storage services are available (see Section 3.9.2).
R118c	<i>Some Platform security services are provided through the identity store.</i>
R118d	<i>Existing privacy services like KODEX can be applied in the same service mesh (see Section 3.9).</i>
R118e	<i>The data integration services are not yet integrated.</i>
R119	<i>This depends on the AI service concept (Section 3.10.6).</i>
R119a	<i>This depends on the AI service concept (Section 3.10.6).</i>
R119b	<i>This depends on the AI service concept (Section 3.10.6).</i>
R119c	<i>This depends on the AI service concept (Section 3.10.6).</i>
R119d	<i>This depends on the AI service concept (Section 3.10.6).</i>
R119e	<i>This depends on the AI service concept (Section 3.10.6).</i>

3.11 Configuration Layer

It is important to recall that all relevant static and runtime information shall be reflected in terms of IVML structures, relations and constraints, while the IVML validation reasoner validates the platform configuration before and at runtime by identifying problems and deviations from validation rules and expected information. The Configuration Layer provides functionality to define applications in terms of the platform IVML configuration on top of the (reusable) services, to dynamically and adaptively optimize the deployment of services and containers and to adapt the use of services at runtime.

Table 22 summarizes all requirements from [13] regarding the configuration. The use of the configuration for resource optimization or adaptation is not listed in Table 22. In this release, we focus on the Configuration component (responsible for the configuration modelling and the instantiation) of the Configuration layer. Optimized container deployment and adaptive operations are deferred to future releases.

Table 22: Specific requirements for the configuration (in addition to the general requirements in Table 2, Table 3)

Requirement	Summary
R8	SPL approaches shall be used for variant management.
R8a	The platform must contain an integrated configuration model for applications, services and platform properties.
R8b	Automated validation of the configuration model
R8c	Automated derivation of platform instances
R12a	The platform can automatically derive the documentation of data processing methods from the configuration model.
R17a	Connectors shall be described in the configuration model.
R19f	The platform shall provide mechanisms for format adaptation or format conversion described in the configuration model.
R19g	The platform shall provide mechanisms for customization or manipulation of metadata as specified in the configuration model.
R20a	Data paths/relations must be defined in the configuration model.
R20b	Data paths/relations can have properties/parameters.
R25c	The platform must manage the available resources.
R31c	The required components to be installed into a container must be specified in the configuration model.
R31b	Containers can contain optional components.
R34	The creation of containers by the platform shall be automated, based on the settings in the configuration model.

Requirement	Summary
R34a	A model validation can be performed before creation or execution to ensure executability.
R34b	The platform can support externally provided containers (e.g., for digital twins).
R36	The platform shall enable configuration settings for resources (read/write).
R40a	RBAC roles can be specified in the configuration model.
R40b	TLS certificates can be specified in the configuration model.
R41a	Directory services must be configured in the configuration model.
R41b	Absence of directory services shall be handled through the configuration.
R42	Further safety mechanisms must be configured uniformly via the configuration model.
R43	Performance targets shall be considered in the configuration model.
R44	The configuration model shall offer IDS-based connectors as optionally configurable.
R64a	The specification of the data fields for anonymization shall be done via the configuration model.
R65a	The specification of the data fields for anonymization of personal data shall be done via the configuration model.
R73e	The data schema for storage services of structured data shall be defined in the configuration model.
R77a	If the platform supports cloud services, the configuration model must offer the use of cloud-based storage services as an option.
R80	Data (including meta-data) shall be described in the configuration model, including data protection classes.
R86	The functionality of the data integration shall be defined by the configuration model.
R89	The platform must allow the data integration write access to data. The data stores shall be defined in the configuration model.
R93	The platform must be systematically configurable in the form of a configuration model.
R94	The platform must support the automatic validation of the configuration model for inconsistencies and errors.
R94a	Validating a configuration model with 50 resources and 5 applications shall be completed in less than 1 second.
R95	The platform must support automatic platform instantiation for a configuration.
R95a	The instantiation of a configuration model with 50 resources and 5 applications shall be completed in less than 15 minutes.
R96	The configuration model must represent optional and alternative platform components/services.
R96a	The configuration model must describe properties of the platform components/services.
R97	The configuration model must include the applications running on the platform.
R97a	An application configuration must contain the configured services for the application.
R97b	An application configuration must contain the configured connectors for the application.
R97c	An application configuration must contain the data paths of the application.
R97d	An application configuration shall contain alternative services.
R97e	The configuration model can allow for application templates.
R98	The configuration model shall support customizations at different times in the software lifecycle.

Requirement	Summary
R99	Information from the configuration model can be made available to other components via internal connectors.
R100	The configuration model can be a decentralized model.
R101	Information provided in the AAS of components/services shall be mapped automatically into the configuration model.
R101a	The transfer of information for a configuration model with 50 resources and 5 applications shall be completed in less than 1 second.
R112a	Parameters of the AI services shall be described in the configuration model.
R112b	Properties of the distribution of AI services shall be described in the configuration model.
R112c	Distribution shall be subject to restrictions for individual AI services.
R113a	Technical dependencies to AI frameworks shall be specified in the configuration model.
R119b	The release of the trained model shall be determined via settings in the configuration model.
R119c	The release of the trained model can be automatic (if specified in the configuration model).
R119e	Changes initiated by a release of an AI model training shall be reversible, e.g., due to configurable criteria.
R120	The configuration model must describe alternative AI components for an AI method.
R122a	The adaptation must be based on the configuration model.
R122c	The adaptation must store its decisions in the configuration model.
R123	Enactment of the adaptation must happen through the configuration.
R123a	The platform can offer mechanisms for updates.
R131	The platform must support the creation and deployment of applications.
R131a	The configuration model must support the specification of applications, their required services, connectors, involved data paths and the needed resources.
R131b	The configuration model must allow for the versioning of applications and services.
R131c	The configuration model can enable the parameterization of applications.
R131d	The configuration model shall support application templates for simplified configuration of requirements.
R131e	The configuration model must describe dependent applications or services.
R131f	The configuration of applications and data paths can be done in a graphical way.
R131g	The platform must support the deployment of applications.
R132a	The configuration model must support application-specific services.
R133a	The platform must know the status of the services.
R133b	The platform must know the status of the running applications.
R134	The platform must support the removal of applications/services.
R134a	The platform must support the removal of deployed applications/services.
R134b	The platform can support the removal of applications from the configuration model.
R135	The platform shall support the update of applications.
R138	The platform can provide an app store.
R139	The platform can provide a user interface.
R141	The platform can enable the visualization of data.

Figure 47 illustrates the design of the configuration component. While the diagram (and the implementation) may appear rather trivial, most of the complexity is in the configuration model, the instantiation process and the underlying framework EASy-Producer.

As already discussed for Figure 2, the configuration model follows the layered architecture of the platform, i.e., each platform layer is represented by a configuration module. Figure 47 just indicates the topmost module, named IIPEcosphere, representing the configuration meta-model, i.e., the configuration options, their structures as well as constraints permitting certain configurations or propagating values among configuration options. We will discuss the model in more details in Section 6. For each platform to be installed, a dedicated platform configuration is created which specifies the AAS settings, the platform data types, the platform services etc. Moreover, for each application a separated (imported) configuration module shall be created, which contains the application-specific data types, the application-specific services as well as the service meshes (directed data flow graphs relating connectors and services) constituting the application. This combined platform configuration is one dedicated instance of IIPEcosphere, in Figure 47 an application configuration taken as input from the Application Layer is illustrated.

The platform instantiation process is defined based on IIPEcosphere meta-model, i.e., an instance of IIPEcosphere can be used as input that defines how the platform shall be instantiated. The platform instantiation process turns the configured information into source code artifacts, setup information, deployment descriptors and executable build scripts. This process also significantly contributes to the invisible complexity of this component. We will discuss also the instantiation process in more details in Section 6.

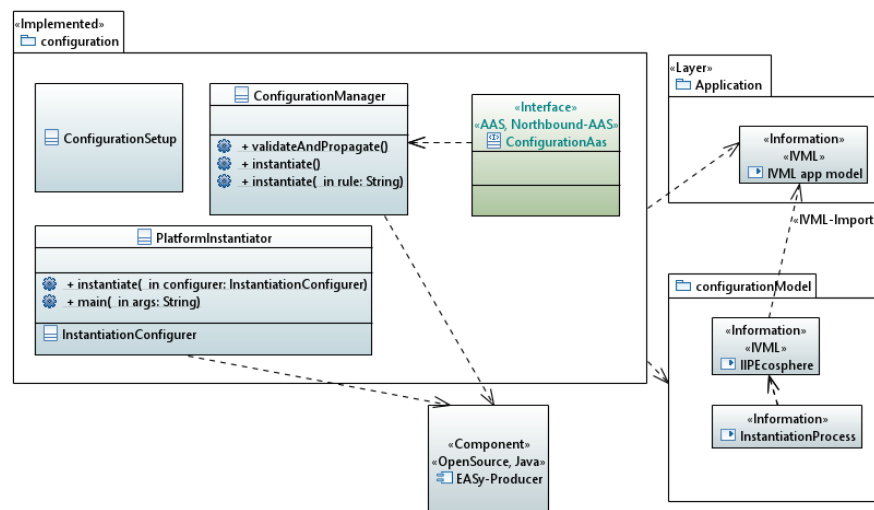


Figure 47: Configuration and instantiation of Definition of applications and orchestration of services (comments cropped)

On top of the models and the instantiation process, the Configuration component just orchestrates the relevant processes. The ConfigurationSetup (read from a Yaml setup file) defines the file system paths where the meta-model, its instance and the instantiation process are defined (meta-model and instantiation process are part of the respective release). The ConfigurationManager ensures the consistency of the operations, currently of loading, validating and instantiating the model. In future releases, also modifications to the actual instance of IIPEcosphere will be provided. The configuration model, the actual settings and in particular the application data flows are reflected in the ConfigurationAas. To ease UI integration, the ConfigurationAas offers operations to read out the data flow graphs in various formats, to write back/delete data flow graphs, to create and delete configuration variables, to change configuration variables and to initiate the application instantiation. For these operations, it is essential that no arbitrary modifications are permitted. The result must always be a valid configuration that can be instantiated. For this purpose, all operations perform a validation through the IVML reasoner and persist the configuration model only if the model is valid. Further, the PlatformInstantiator realizes a command line tool to perform the basic operations

of the ConfigurationManager, i.e., to allow a user to instantiate the platform and the defined applications. The PlatformInstantiator offers various modes, ranging over the instantiation of interfaces for applications, the full instantiation of applications, the instantiation of platform components, etc.

The configuration model and the platform instantiation are subject to regression testing in the continuous integration. While the validation within the configuration model is currently not as extensive as it could be, partial results of the platform instantiation have been validated for functionality, e.g., the instantiation of executable platform components, the creation of several stream-based IoT example applications as well as the automated low-code generation of connectors. In particular, the connectors were validated in the platform use case studies with partners (cf. Section 3.6.3.3), i.e., the respective input/output formats have been modeled in configuration, the connector service integrations have been generated and the intake has been validated. To demonstrate the setup of the platform, the platform instantiation as well as the creation of example service artifacts is part of the Docker platform containers provided on DockerHub.

Following the structure of the previous section, we discuss now the implementation of the configuration requirements. However, so far, we did not detail the structure of the IVML model and the capabilities of the instantiation. This information is provided in Section 6, because we focus here on the architectural side. To avoid two separate discussions of the realized requirements, Table 23 includes forward pointers to Section 6 and summarizes already the requirements state explained there.

Table 23: Review of realized⁸⁰ requirements for the configuration (based on Table 2, Table 3 and Table 22)

Requirement	Summary
R8	SPL approaches are used for variant management.
R8a	The platform contains an integrated configuration model for applications, services and platform properties.
R8b	Automated validation of the configuration model is supported in terms of the constraints in the variability model and the EASy-producer IVML reasoner (cf. Section 6).
R8c	Automated derivation of platform instances is supported through the instantiation process (cf. Section 6).
R12a	<i>Derivation of the documentation of data processing is currently not supported.</i>
R17a	Connectors are part of the configuration model (cf. Section 6).
R19f	The configuration model allows specifying data mapping/format conversion functions for connectors, <i>not yet for all services</i> . Data serializers are implicitly derived during instantiation (cf. Section 6).
R19g	<i>Currently the configuration model does not support mechanisms for customization or manipulation of metadata.</i>
R20a	Data paths/relations are defined in the configuration model in terms of service meshes (cf. Section 6).
R20b	Data paths/relations can have properties/parameters although currently only the name is specified (cf. Section 6).
R25c	<i>The platform must manage the available resources. Resources are foreseen through the conceptual foundation of the meta-model but not used (cf. Section 6).</i>
R31c	<i>The required components to be installed into a container are currently not linked to resources.</i>
R31b	<i>Containers can contain optional components. Currently, services and resources are not linked.</i>
R34	The automated creation of containers is supported by the platform.

Requirement	Summary
R34a	A model validation is performed before instantiation and through the ConfigurationManager can be performed before further platform operations.
R34b	<i>Currently no externally provided containers are supported.</i>
R36	<i>The platform shall enable configuration settings for resources (read/write). These resource configuration settings may be reflected by the device management into the configuration.</i>
R40a	<i>RBAC roles are currently not specified in the configuration model.</i>
R40b	<i>TLS certificates are currently not specified in the configuration model.</i>
R41	Security mechanisms shall be integrated with common directory services.
R41a	Directory services <i>can be integrated</i> through the identity store interface.
R41b	<i>The absence of directory services is currently not considered by the configuration.</i>
R42	<i>Further safety mechanisms are currently not part of the configuration model.</i>
R43	<i>Performance targets are currently not part of the configuration model.</i>
R44	<i>The configuration currently does not offer IDS-based connectors as optionally configurable.</i>
R64a	<i>The specification of the data fields for anonymization is currently not supported by the configuration model.</i>
R65a	<i>The specification of the data fields for anonymization of personal data is currently not supported by configuration model.</i>
R73e	<i>The data schema for storage services of structured data is currently not linked to data storages.</i>
R77a	<i>The platform currently does not support cloud services.</i>
R80	Data (including meta-data) shall be described in the configuration model, including data protection classes. Data types are supported, <i>meta data or protection classes are currently not part of the configuration model.</i>
R86	<i>The functionality of the data integration is currently not part of the configuration model.</i>
R89	<i>Currently no data stores are defined in the configuration model.</i>
R93	The platform is systematically configurable through a configuration model.
R94	The platform does support the automatic validation of the configuration model.
R94a	Validating a configuration model with 50 resources and 5 applications shall be completed in less than 1 second. <i>The models currently do not reach this size, but on example configurations validation is currently not a bottleneck.</i>
R95	The configuration model does support automatic platform instantiation.
R95a	The instantiation of a configuration model with 50 resources and 5 applications shall be completed in less than 15 minutes. <i>So far, test configurations up to 3 applications including all platform components require less than 3 minutes on a computer in the same network as the (snapshot) code repository.</i>
R96	The configuration model includes optional and alternative platform components/services (cf. Section 6).
R96a	The configuration describes properties of the platform components/services (cf. Section 6).
R97	The configuration model defines applications running on the platform (cf. Section 6).
R97a	An application configuration contains the configured services for the applications (cf. Section 6).
R97b	An application configuration contains the configured connectors for an application (cf. Section 6).
R97c	An application configuration must contain the data paths/relations of an application (cf. Section 6).

Requirement	Summary
R97d	An application configuration does allow for alternative services (via families, cf. Section 6).
R97e	<i>The configuration model currently does not allow for application templates.</i>
R98	The configuration model does allow for customizations at different times in the software lifecycle <i>although not all relevant ones are defined</i> (cf. Section 6).
R99	<i>Information from the configuration model is currently not made available to other components via internal connectors.</i>
R100	<i>The configuration model is currently a centralized model.</i>
R101	<i>Information provided in the AAS of components/services is currently not mapped automatically into the configuration model.</i>
R101a	<i>The transfer time for a configuration model with 50 resources and 5 applications is currently unknown as R101 is not realized.</i>
R112a	<i>Parameters of (AI) services are currently not described in the configuration model.</i>
R112b	Properties of the distribution of AI services shall be described in the configuration model. <i>Currently, the configuration contains only the information whether a service is distributable.</i>
R112c	<i>Currently no distribution applies.</i>
R113a	Technical dependencies to AI frameworks are available in the configuration model. <i>The automated instantiation into containers is still in realization.</i>
R119b	<i>The release of a trained model is currently not considered in the configuration model.</i>
R119c	<i>The release of the trained model is currently not considered in the configuration model.</i>
R119e	<i>Changes initiated by AI model training are currently not subject to configurable release or quality criteria.</i>
R120	The configuration model must describe alternative AI components for an AI method. This is realized in conjunction with R97d.
R122a	<i>Currently, runtime parameters and service families can be specified in the configuration model and are considered during code generation.</i>
R122c	<i>The adaptation is not part of this release.</i>
R123	Enactment of parameter happens through the configuration/code generation, enactment of alternative services through <i>data plugins</i> .
R123a	The platform offers interfaces for updates.
R131	The platform allows for the creation and deployment of applications.
R131a	The configuration model supports the specification of applications, their required services, connectors and involved data paths.
R131b	The configuration model allows for the versioning of applications and services.
R131c	The configuration model allows for the the parameterization of <i>services</i> .
R131d	<i>The configuration model currently does not support application templates for simplified configuration of requirements.</i>
R131e	<i>The configuration model currently does not describe dependent applications, but service chains in service meshes.</i>
R131f	<i>The configuration of applications and data paths is currently not done in a graphical way as no UI is provided.</i>
R131g	The platform supports the deployment of applications.
R132a	The configuration model does support application-specific services.
R133a	The platform must know the status of the services. <i>Currently no runtime data is reflected in the configuration.</i>
R133b	The platform must know the status of the running applications. <i>Currently no runtime data is reflected in the configuration.</i>
R134	The platform supports the removal of applications/services.

Requirement	Summary
R134a	The platform supports the removal of deployed applications/services.
R134b	The platform supports the removal of applications from the configuration model via the configuration AAS.
R135	The platform shall support the update of applications. <i>The configuration layer currently does not provide detailed configuration manipulation operations.</i>
R138	<i>The platform provides mechanisms to realize an app store, in particular the dynamic extensibility of the configuration model (although not yet through the configuration AAS).</i>
R139	<i>The platform provides a user interface, which is in development.</i>
R141	The platform provides a visualization of its collected monitoring data, <i>but not yet the processed data.</i>

We conclude, that basic requirements for this layer are implemented, in particular also for services and applications ([13] only states “the application”). However, there are many (cross-cutting) requirements for the configuration in [13] and in several cases the underlying platform components are not realized so that configuration modeling for those requirements is useless at the moment. These components/requirements are scheduled for future releases.

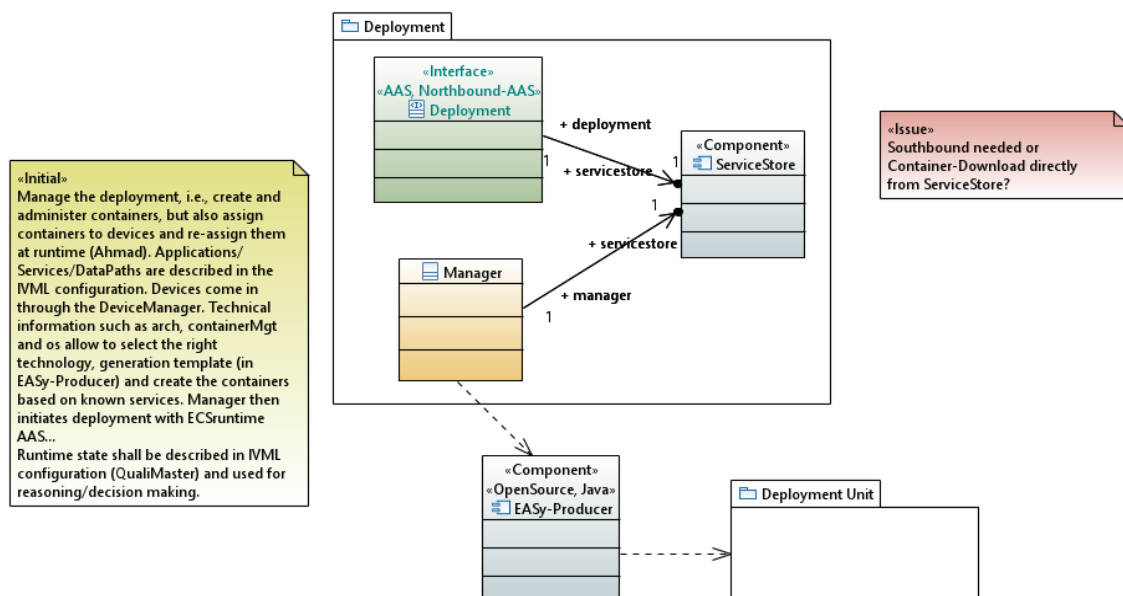


Figure 48: Optimization and adaptation of service/container deployment




Figure 49: Adaptation of running services (comments cropped)

3.12 Application Layer

Ultimately, the Application Layer represents individual applications, i.e., it is the actual home of the application configurations to be installed, the generated artifacts and additional application-specific (handcrafted) components and services. The overall picture is depicted in Figure 50.

Currently, this layer does not really exist as platform instance/application configurations are defined as part of the tests of the Configuration Layer or on the command line of the respective tooling. Thus, generated and packaged artifacts are currently belonging to the Configuration Component (temporary, generated artifacts folder). The setup and the application layer will change in the next releases.

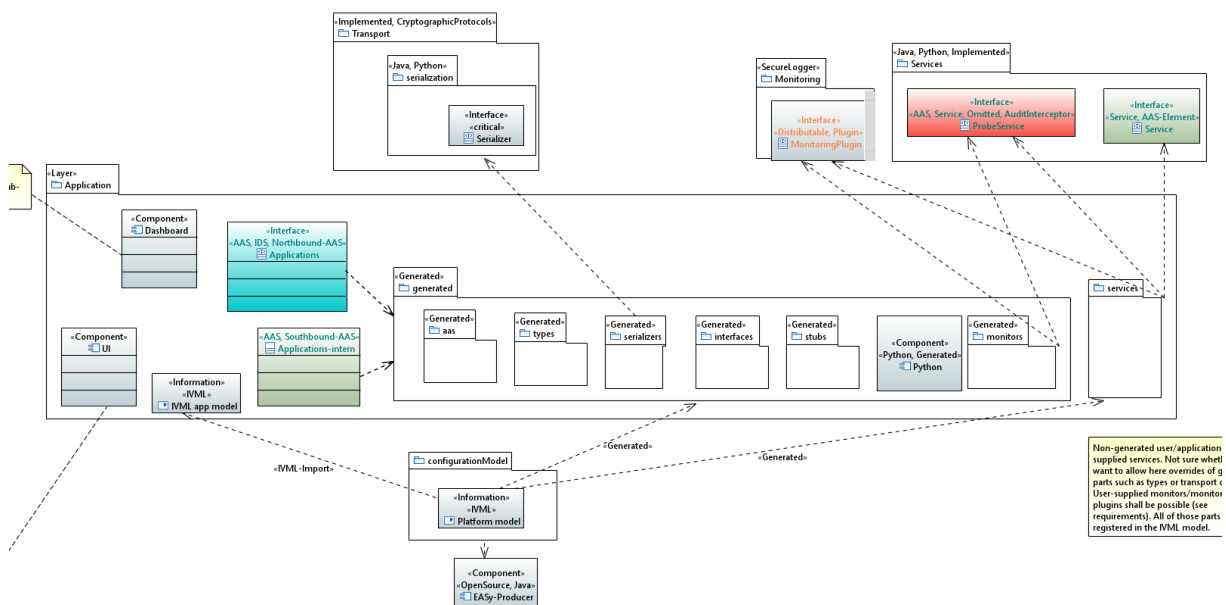


Figure 50: Application Layer (comments cropped)

3.13 Platform Server(s)

As discussed above, the IIP-Ecosphere platform consists of several layers and many components. However, so far there also is a component that provides the setup and lifecycle mechanisms for the central IT-side of the platform, e.g., powering up the platform AAS service. At a glance, this component does not provide new functionality or concepts and may not be worth mentioning. In fact, it is a vital

part for later platform instantiation, as it defines how central services can be configured, instantiate and how these services are started. Moreover, it provides an initial simple command line interface to operate with the IIP-Ecosphere platform, e.g., to start containers or services.

Figure 51 depicts the structural design of the platform component through using the server implementations and server-related parts defined in all layers and components discussed before. As stated above in this chapter, this component serves for two purposes:

1. Powering up the servers to run the IIP-Ecosphere platform. Therefore, the component defines a lifecycle descriptor (`PlatformLifecycleDescriptor`), which reads information from the `PlatformSetup` representing the YAML setup file. The lifecycle descriptor is loaded via JSL into the `LifecycleHandler`, which, in turn, is called by the platform component during its main program. During this startup process, all “installed” lifecycle descriptors (e.g., the descriptor for the network manager; the platform instantiation is responsible for this) are also started up. As part of the startup also the platform AAS is constructed, which contains the platform “nameplate” (`TechnicalInformation` sub-model [2]), further software-specific information (`Platform` sub-model) as well as a listing of all available application `Artifacts` (service artifacts, containers, deployment plans).
2. Providing a simple command line interface (`Cli`) to experience the operations of the IIP-Ecosphere platform. For the command line interface, simple means that we (currently) provide access to some, particularly rather low-level functionality of the platform, which incrementally shall be taken over into the platform, e.g., which service shall be deployed where shall be part of the deployment component in the configuration layer. In addition, first higher-level commands such as executing a deployment/undeployment plan are provided, which adds/removes services implementation artifacts and starts/stops contained services that shall be distributed across multiple devices. The command line interface does not rely on the lifecycle mechanism, but on the `PlatformSetup` and, in particular, on the AAS clients of the service and the resources layer to ease executing the operations defined there. Figure 52 illustrates an example interaction with the interactive mode of the command line interface, here turning into the resources commands, showing the commands for resources (`help`), listing the available resources and, finally, ending the client. For the single resource shown in Figure 52, in particular the integrated container manager (for Docker) and various initial runtime measurements for disk and memory allocation are shown. It is important to emphasize that the command line performs its operations via the platform AAS and the respective AAS clients for services and the ECS runtime.

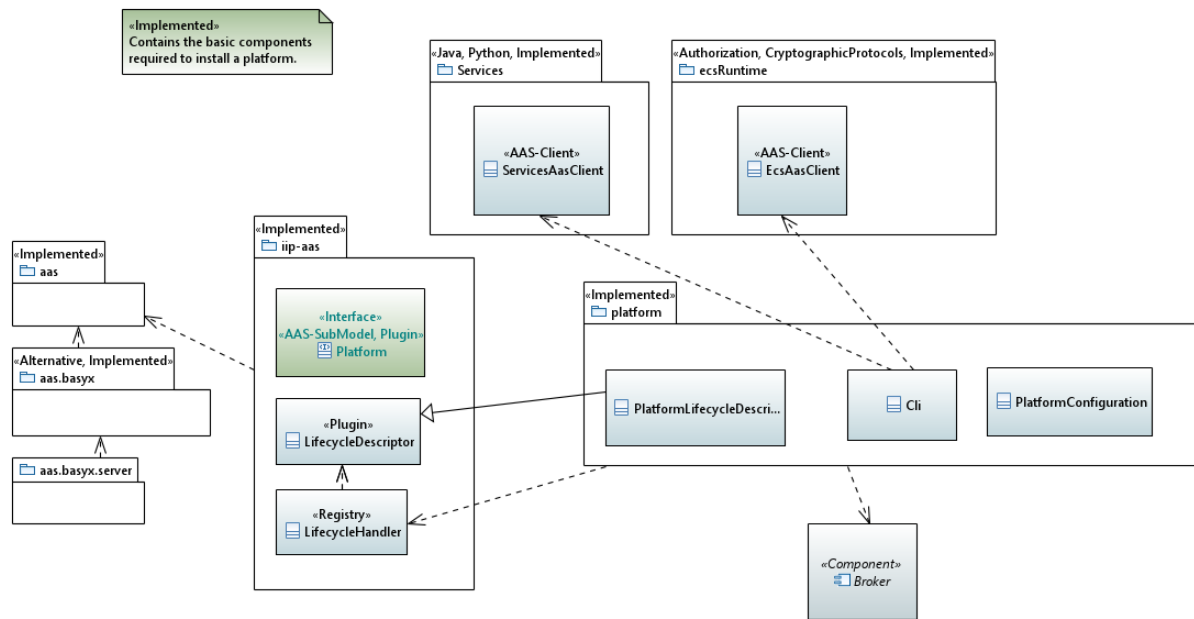


Figure 51: Platform server(s) component

```

IIP-Ecosphere, interactive platform command line
AAS server: http://127.0.0.1:9001
AAS registry: http://127.0.0.1:9002/registry
Type "help" for help.
> resources
resources> help
  list
  help
  back
resources> list
- Resource a005056C00008
  systemdisktotal: 1023887356
  systemmemorytotal: 2147483647
  simplemeterlist: ["system.cpu.count", "system.cpu.usage",
    "system.disk.free", "system.memory.free"...]
  containerSystemName: Docker
  systemmemoryfree: 2147483647
  systemdiskfree: 464061712
  systemmemoryused: 2147483647
  systemdiskusable: 464061712
  systemmemoryusage: 0.5555296172875698
  systemdiskused: 559825644
resources> back
> exit

```

Figure 52: Interaction with the preliminary interactive platform command line interface.

Using the platform command line interface, we validated the interaction among the components. Therefore, we started platform, ECS runtime and service manager component as individual programs. Through the command line interface, we validated the resource represented by the ECS runtime and started a simple generated application (cf. Section 6). We identified here the following issues:

- BaSyx issues exceptions when checking whether a non-existing AAS exists through trying to access it.

- Long running commands such as starting services are currently rather quiet on the command line interface, i.e., they do not show intermediary steps while the logs on the respective device indicate the actual state. AAS do not support return streams, so either polling from the caller or transmitting the results via the Transport Layer could be options for improvement.

We also validated the execution of services in a service manager container, starting and stopping of containers via the platform and the ECS runtime execution in terms of a (Docker-out-of-Docker) container. Please refer to Section 8.4 on how to install, instantiate and containerize the IIP-Ecosphere platform, i.e., to perform the steps that we also executed for validating the command line interface and the instantiated platform components. The platform CLI also supports creating snapshots of the platform AAS¹⁰⁴ that can be explored with the AASX Package Explorer¹⁰⁵.

As the platform layer shall also be used as basis for a Management User Interface (cf. Section 3.14), additional information that is usable on that layer may be required. For this purpose, we added a listing of available artifacts (containers, service artifacts, deployment plans) to the Platform AAS, so that they can be selected, inspected or executed on UI level. Akin, further information and interfaces regarding the configuration will be required.

3.14 Management User Interface

As already stated in [41, 13], no real user interface was scheduled for the IIP-Ecosphere platform in the grant agreement. However, the value of an accessible and usable user interface for a platform (over a simple CLI as discussed in Section 3.13) is evident. For this reason, the platform team of IIP-Ecosphere aims at providing a management user interface, i.e., a web user interface that allows for managing the platform operations such as starting or composing an application.

Currently, the user interface is initial and focuses on displaying the information that is accessible through the CLI as well as basic management operations provided by the CLI. Here, the Platform AAS forms the information model the Management UI can rely on and one interesting question is whether it is possible to realize an efficient Management UI based on an AAS. Challenging tasks include reading the nested AAS REST structures or calling AAS operations, as BaSyx does so far not provide TypeScript support. While for some operations on a command line we can assume that the user can look them up, e.g., the URI to a service artifact, this may not be the right approach for a Web UI. Thus, as part of creating a management User Interface, we also have to provide additional information, e.g., on the available artifacts, containers or deployment plans and to make them “executable” via the UI.

The user interface requires some form of setup, in particular knowledge about the installation location of the platform AAS servers. To resemble UI release versions and integration with the platform instantiation and installation approach, the management UI allows for compiling the TypeScript code for Angular while allowing for an external setup through a JSON file¹⁰⁶. The compilation happens as part of the Continuous Integration of the IIP-Ecosphere platform, the customization during the instantiation based on information in the configuration model. The platform instantiation turns the compiled Management UI into an instantiated version, where in particular the settings in the Angular environment are adjusted or respective start scripts, e.g., for an Express webserver are generated.

¹⁰⁴ Currently it seems that BaSyx allows only a single snapshot per run. This may change in future versions.

¹⁰⁵ <https://www.plattform-i40.de/IP/Redaktion/DE/Newsletter/2019/Ausgabe21/2019-21-Praxisbeispiel2.html> tested with version 2021-08-17.alpha.

¹⁰⁶ <https://mokkapps.de/blog/how-to-build-an-angular-app-once-and-deploy-it-to-multiple-environments/>

3.15 Test support

So far, we focused on the elements to construct the platform, the services as well as the applications. In this section, we provide a cross-cutting overview on the testing support, in particular to answer the question, how the platform supports the user in testing his/her own services and applications.

Besides internal component and service testing, in particular of services supplied with the platform, it is essential to test user-developed services as well as their interactions in an application. At a glance, the answer might be to construct a unit test and to test the supplied code. However, reality is not so trivial, as, e.g., connectors may be based on external devices or their server instances, e.g., MQTT broker or OPC UA server, and these devices may not be available in certain testing situations. Moreover, setting up a test for a single spring-based service involving a Java or, in the more complex case, a Python service (with involved Java integration and Python service environment) requires much internal knowledge and may even lead to problems if the service layer is equipped with alternative service execution engines that are not considered by the test.

As a general advice, we recommend to apply testing to all levels of an application, ranging from tests of the code that you supply up to entire applications. In particular, testing of services before running the application instantiation and packaging process can usually save much time, i.e., as usual, getting rid of failures in individual services shall be performed before integration tests of the application.

Thus, the platform offers different forms of testing support that we will summarize here:

- **Testing an individual connector:** Although we demonstrate in the RoutingTest regression test how to write and integrate a self-supplied simple connector, typically, a connector is a kind of platform-supported service. Thus, the connector itself, either a model-based or channel-based connector is already tested sufficiently by platform tests. The instantiation process wraps the connector into a service and adds generated input/output data translators or user-supplied data translators or event handlers based on the data specification of the configuration model to turn the generic connector type into an application-specific connector. Although the user-supplied parts shall be tested individually, it always remains unclear whether they will work correctly in the interaction context of a connector. For this purpose, the generated code for the configured connector sets up an environment that the connector can be executed individually. By default, the connector runs against the configured device, e.g., the OPC UA connector against the OPC UA server of the configured device. As this requires the device at hands, the test may fail in a CI environment. Thus, it is possible to execute the test as a stand-alone program (for the target environment) and to mock the connector (as set up in the configuration model). When the connector is mocked, the code generator wraps the connector differently into a service, i.e., it creates a connector instance, but detangles it from the execution while pretending to feed the connector output with data from a JSON file (via functionality provided by the service environment). Such a mocked connector allows you to experiment with different data settings even in a CI environment.
- **Testing individual services:** For a service it is more likely to integrate self-supplied code. This code shall be tested individually. However, as for a connector, such tests lack the full environment, e.g., the surrounding Python environment or the service lifecycle. Here, the generated test cases set up an environment that fits to the actual service execution environment, per default Spring Cloud stream. Akin to mocked connectors, a JSON file determines the input data (which may consist of multiple data type instances, and may define a timed ingestion behavior), which is fed through the DataWrapper into the service through the actual mechanisms of the service environment or the platform, e.g., the data transfer. Resulting data, synchronous or asynchronous, is received by the test and basically emitted. In the generated version, all data received from the service under testing are asserted as true as

we do not employ a data correctness specification at the moment. Basically, the test only ensures that data shall come out of a test (assuming that data was fed into it). Statistics about received data types are collected by default and can be used for simple asserts. Further, you may extend the test to determine more complex assertion behavior and to turn it into a real regression test.

- **Mocking of applications:** Although components may be working after applying the test opportunities discussed above, there is no guarantee that the integrated application will be working. Here, again, mocking may be required as not all devices or even software environments are available, that you would need for mocking tests. Connectors can be mocked as above, e.g., if certain devices and their OPC UA servers are not available. Services can also be mocked, by replacing the service classes defined in the model, e.g., by extensions of the services that disable some functionality. One example is to replace a TensorFlow-based Python AI script (with one reason that you may not have required GPU capabilities at hands) by a simple mock script and to tell the generation to use the mock script instead of the original script. To ease this, the IVML configuration model allows you to define additional variables, e.g., one for testing and to define functions that return the actual or, in testing, the mocking service class. Testing may be enabled, e.g., by modifying the testing variable before instantiation or by defining two different top-level models and to change the model to use in the calls to the PlatformInstantiator tool.
- **Testing the application:** Despite all tests, ultimately also the application with all services and all required devices in place must be tested.

It is important to mention that only testing on application-level may include all resources and service implementations in the final form as it will be deployed. Thus, accidental overlaps of resources, e.g., identity stores may only be detected when running an integrated application.

The requirements documents [41, 13] even demand in-place pre-deployment tests. Currently, the platform does not offer functionality for these optional (but important) requirements.

4 Architectural Constraints

Besides structure and communication sequences, often an architecture explicitly or implicitly defines constraints that must be obeyed by an implementation. We summarize and explain the constraints for the IIP-Ecosphere platform here:

- C1. Higher layers and contained components are allowed to have **dependencies** only to downstream layers and components, if possible only to the directly adjacent lower layer. This constraint is induced by the basic layered architecture style of the IIP-Ecosphere platform.
- C2. As an exception from C1, the **ECS runtime shall not depend on the Services Layer** so that the services layer can be installed separately (as explained in Section 3). Both, Services Layer and ECS runtime may depend on certain classes of the services environment.
- C3. Wrapped **singleton components** or libraries shall not be called by other components than the wrapper itself. Basically, this applies to transport and connector protocols, the AAS implementation (BaSyx), but also for container management libraries such as Docker. This constraint intentionally focuses on singleton components/libraries, as some libraries may occur in multiple component dependencies, e.g., the stream processing framework due to the need for different protocol/binder implementations. In turn, this also applies to some transport/connector protocol client implementations. Another exception is `support.aas.basyx.server`, which is allowed to access (as the only component) `support.aas.basyx` as it represents the server component with full dependencies.
- C4. **Support components** for C3 shall be realized as **optional components**, e.g., the Spring service environment refining the generic Java environment. There shall be no references into such components except for refining components. In particular, generic components shall not reference their specialized components. For providing access to the specialized implementation, descriptors, factories or facades are to be used where the implementation is provided by JSL.
- C5. **Protocol servers** for testing such as Apache Qpid, HiveMq or Moquette shall be in testing components and no other component shall directly use classes from them (although Maven requires explicitly naming also those transitive dependencies). These testing servers may be used during platform instantiation to provide a broker/server for a selected protocol.
- C6. Production code must **not have dependencies to alternative or optional components**. As a rule of thumb, generic components without “suffix” names (representing the generic part of a component) shall not directly access related optional/alternative components indicated by “suffix” names, e.g., `transport` is the generic transport layer while `transport.amp` the alternative for the AMPQ protocol. This applies to the support layer (no access to `BaSyx/Server`), the transport binders (e.g., `transport.spring` vs. `transport.spring.amqp`), the connectors, the service environment (`services.environment` vs. `services.environment.spring`), the services (service manager `services` vs. `services.spring`) the ECS-Runtime (e.g., `ecsRuntime` vs. `ecsRuntime.docker`) etc. In contrast to production code, test code (Maven scope “test”) may declare dependencies to specific alternatives to allow for functional testing, e.g., to rely explicitly on the AMPQ transport protocol. Although alternatives shall be tested equally on their level, it is also clear that component testing with specific alternatives just shows the functionality for the assumed/selected alternatives.
- C7. **Generated artifacts** shall be separated from manual code (usually an own top-level folder such as `gen`) and generated artifacts shall not be modified as they may/will be re-generated upon request.
- C8. **Implementation of services** shall be separated per service, so that services can be composed/integrated free of other dependencies. For convenience, in testing code, we may

intentionally validate this rule, e.g., `test.configuration.configuration` implements all service artifacts for all tests in `configuration.configuration`.

- C9. **Exception handling** is often not considered a topic for architectural constraints. However, the basic decisions on how and where to use/handle exceptions are important as they enforce certain responsibilities. Moreover, some architecture modeling languages like UML allow for the specification of exceptions. Exceptions indicate abnormal situations in the program execution that shall not be handled by normal program code rather than by stopping the execution at the point of occurrence and tracing back the method calls until the exception at hands is handled (or on top-level it terminates the program or the actual thread). While often programmers try to handle an exception at the point where it obviously occurs (in Java, where a checked exception could be thrown that must be handled), we believe that in most cases the caller, i.e., the cause of executing the code that throws the exception shall be informed, which does not mean that each exception must be transported and handled in top-level code. For example, consider some complex data format processing code, e.g., reading an AASX file for an asset administration shell. If we handle an Input-Output exception in that code, the caller does not know that and why the format processing fails. Let us now assume, that reading the AASX file was triggered by the ECS runtime when building the AAS of the ECS runtime, e.g., to link device vendor and ECS AAS. Here, the lifecycle handler of the ECS runtime (more or less top-level code) that starts the creation of the ECS AAS is not interested in why an AASX file processing fails. However, the code creating the AAS trying to establish the AAS link is better suited to handle the exception, e.g., to insert an empty link or to log the problem. Thereby, logging (cf. Sections 2 and 8.1) is often not the right answer to an exception, in particular not emitting an exception stack trace to the console (which may not be logged properly). In contrast, the programmer shall think about handling the exception in a manner that processing can succeed, e.g., inserting an empty link into the AAS rather than no AAS property at all, which may cause failures in other parts of the system relying on the assumption that such a property exists. In particular the type of the used exceptions shall be selected carefully (cf. Sections 8.1).
- C10. Apply **defensive logging**, i.e., carefully think about what is an “error”, a “warning”, an “information”. Errors shall only be emitted if a component will fail to operate. If the component can compensate this, e.g., by falling back to some strategy or default plugin, then a warning is more adequate.
- C11. **Logging setup/filtering is decided during integration, not before.** As some “bigger” components like BaSyx, Spring or even Apache QPID-J ship with their own ideas how to set up and configure logging, deferring the logging decision (SLF4J, LogBack, Apache, etc.) and the setup what to log as long as possible. Thus, all components must not define a concrete logging implementation in their production code dependencies, only in their test dependencies. Ultimately, the code generation that is performing the integration automatically must know whether further dependencies are needed, existing dependencies can be used and how the logging of the components at hands shall be set up.

It would be desirable to check and enforce these dependencies. However, so far tools that we tried, e.g., in the continuous integration, failed for multiple components using a central or even adequately distributed rule set as they require an application rather than a component to be checked. We will try to find and integrate a feasible tool as soon as possible.

5 Asset Administration Shells

As stated above, the IIP-Ecosphere platform heavily relies on asset administration shells (AAS) to describe the capabilities and interfaces of its components. Currently, only few standard structures for AAS/sub-models exist while many are still in development, e.g., a software type-plate [2, 43] or a description of qualities of service (QoS). However, it is not feasible for the work on the IIP-Ecosphere platform to wait until such standards are defined. Thus, we follow an agile and pragmatic approach to AAS modeling here:

1. As long as **no guidelines** for AAS in IIP-Ecosphere exist, the IIP-Ecosphere platform will draft AAS that contain the most relevant information for its operation, i.e., for now the IIP-Ecosphere platform relies on AAS prototypes. All names and sub-structures shall be defined in terms of constants so that names and structures can be adapted (within limits). AAS of the IIP-Ecosphere platform shall be tested individually and also in integration settings to handle and to judge the impact of modifications. Tests shall also rely on the defined constants rather than on local String literals. In this stage, we pragmatically focus on AAS describing instances.
2. Discussions with third parties on **(proto-)standardized AAS** structures are ongoing. As soon as results from these discussions are available, a guideline for AAS modeling in IIP-Ecosphere shall be drafted. The experience made with prototyping AAS in the IIP-Ecosphere platform will be considered in these guidelines. At this point, also type AAS shall be provided.
3. The IIP-Ecosphere platform will modify the AAS prototypes and augment the information (e.g. ECLASS references) to **comply with the guidelines**. This may lead to a re-structuring of the AAS prototypes.

With this approach in mind, we designed and partially realized the prototypical IIP-Ecosphere AAS structure shown in Figure 53. As already explained in Section 3.1.2, we separate between AAS describing an (external) artifact and internal information (usually in sub-models). AAS do exist for

- The platform AAS with its various sub-models like name plate, dynamic network port assignment, transport setup, (S3) storage access, (available) artifacts such as containers or deployment plans, installed connector/service types and their utilized data types, the device management and available devices (with installed/running containers, installed service artifacts, running services).
- Further assets represented in their own AAS like devices, service¹⁰⁷ or composed applications (with vendor information). Device and service AAS are linked from the respective platform submodels to make the information in the AAS available. For each application running on top of the platform, an AAS shall be provided (currently via the TraceToAasService discussed in Section 3.7.3.1), which states the creator of the App but lists also the utilized services and may provide application specific operations.

For the platform AAS and its sub-models, we distinguish between installed/available descriptors and their active instances at runtime, in particular as in many cases only the active instances provide the full information about in/outgoing types. Examples are in particular the connectors, the services and their relations, the containers etc. These structures are dynamic, i.e., they change due to installed components as well as due to instantiated/terminated instances. This is in particular the case for connectors and services, subsequently also for applications. Some sub-models are active, in particular

¹⁰⁷ So far, the software nameplate is not available as standard. As platform-provided services typically involve two vendors, the actual service creator and the organization that created the integration into the platform, a linking of AAS would be required. So far, this linking is not realized, rather than just one AAS is provided.

those providing operations. One example for an active AAS is the optional netMgt submodel, which provides access to the local/global NetworkManagement defined in the Support Layer.

It is important to emphasize that the structure shown here is not static. It is dynamic in its elements as explained above, but it is also dynamic in its overall structure and contributions, in particular if the AAS is centrally deployed and parts are added remotely. A specific example is the relation between resources and services. When an ECS runtime comes up, it contributes itself to the resources collection. When a service manager starts, it contributes further operations to the resource it is running on, i.e., both Layers contribute into the same AAS sub-model (elements collection), because in this case the components have information and operations that they only can share individually but that are part of the same topic, namely the runtime interface of a resource.

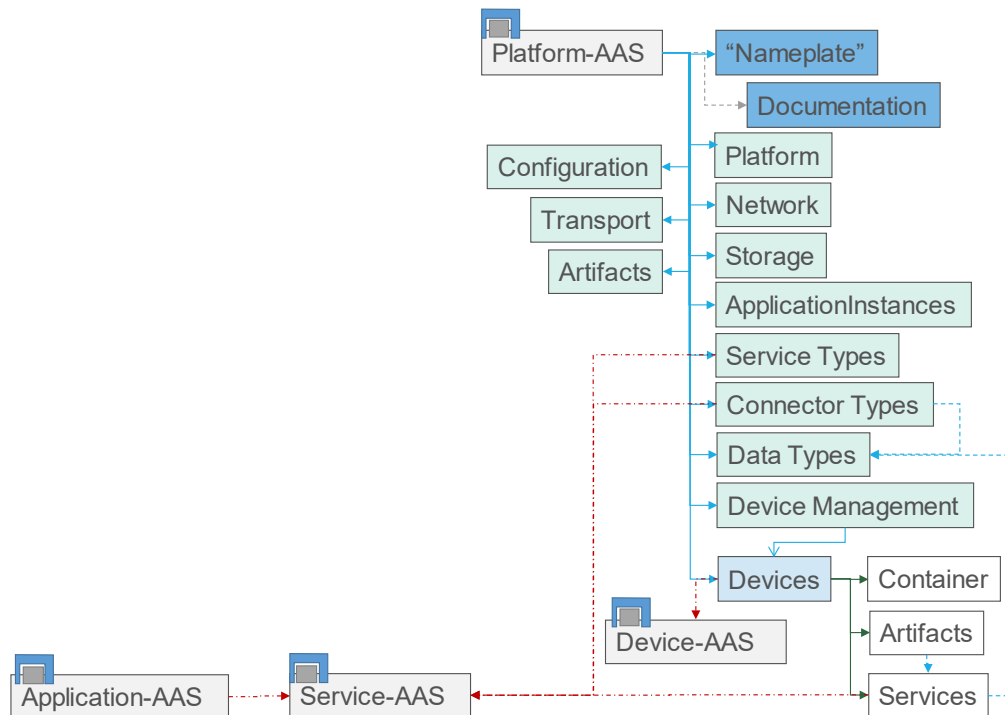


Figure 53: AAS structure of the IIP-Ecosphere platform (preliminary, incomplete)

Figure 54 depicts a screenshot illustrating a fragment of the IIP-Ecosphere platform AAS in the AASX Package Explorer, i.e., an excerpt of the full AAS shown in Figure 53.

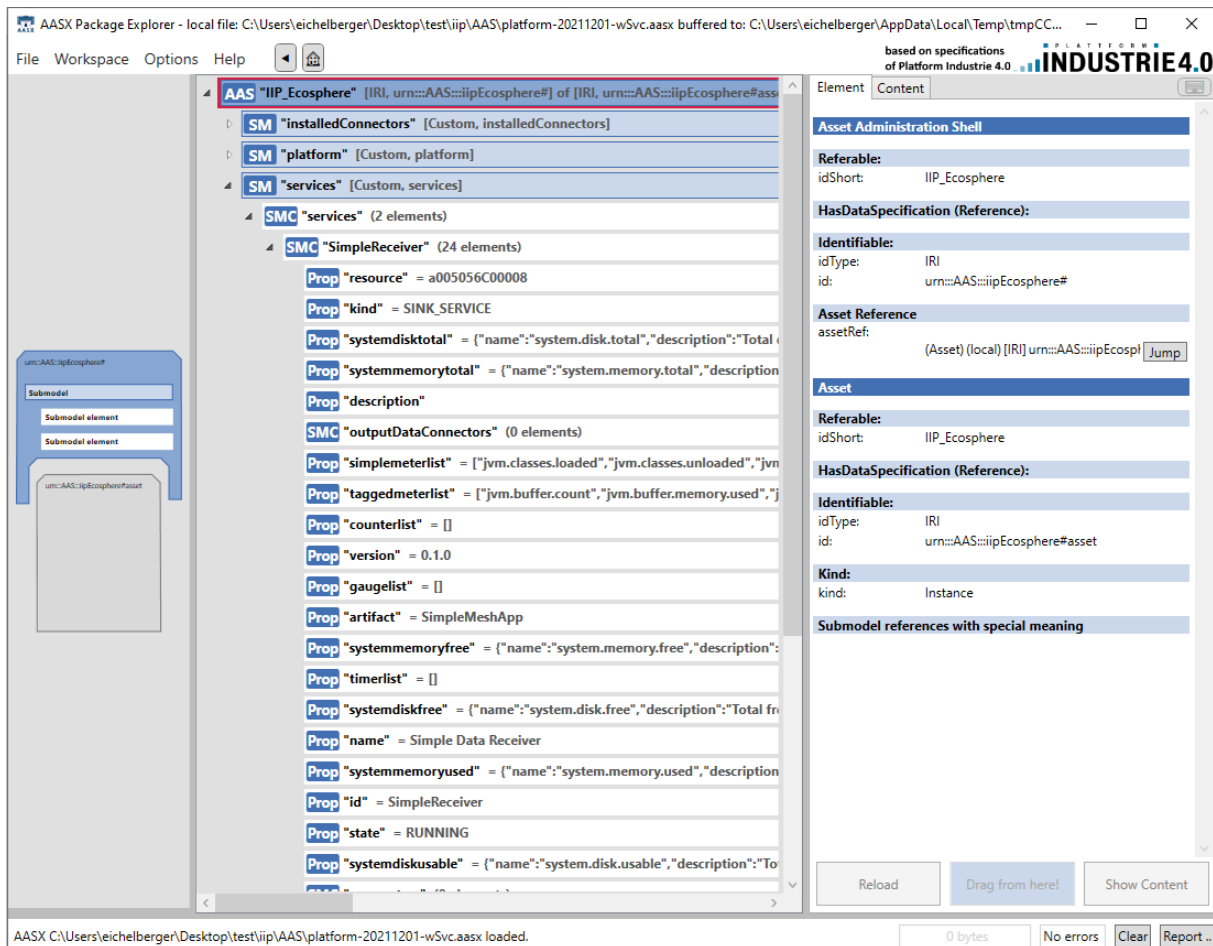


Figure 54: IIP-Ecosphere AAS in the AASX Package Explorer showing a running service (*SimpleReceiver*).

As the IIP-Ecosphere AAS is rather dynamic, we can already draw some conclusions on lessons learned with BaSyx (based on the integrated version through the support layer):

- Remotely deployed AAS with operations and properties realized in terms of attached functors typically require uniquely serializable functor objects, i.e., they do not work with simple lambda functors or serializable lambda functors.
- When obtaining a remotely deployed AAS, the AAS is turned into a serialized format as already briefly mentioned in Section 3.7.3, i.e., all functors such as getters, setters or operations are serialized, to obtain the values of the properties the getters are even executed. If getters are bound to an AAS implementation server, that server must be ready to serve connections at the point in time when the remote AAS is requested (which may happen in parallel initiated by other components) and currently for each property a network connection is created by the respective BaSyx connector and the value is requested. This seriously affects the performance of obtaining and using a remote AAS. It happened to us that in such a situation a potentially endless loop occurred forcing us to re-think a rather obvious implementation approach in terms of getter functors. As discussed in Section 3.7.3, we suggest using functors that map to local data rather than to remote data. The local data object may be updated in parallel through a different process, e.g., a Transport Layer connector. Dependent on the implementation, each serialized AAS then has its own remote data object, leading to a distributed setup of AAS that can be kept up to date via Transport Layer mechanisms. Directly writing values into an AAS might be an alternative, but in the remote deployment case, the serialized AAS implicitly performs update requests on the original remote AAS, i.e., probably leading to reduced performance.

- When writing larger portions of structured data, in particular binary data, there is a conversion problem in the BaSyx version that we are using. Types like Base64Binary are not handled correctly. Currently, we encode such data through a Base64 String encoder (similar to the contents of File Data Elements).
- The IIP-Ecosphere abstraction appears to be easier to use and requires less code than plain BaSyx [3], but this was a design goal. Moreover, the AAS implementation can be replaced seamlessly, also by a non-AAS interface realization.
- So far (as far as we know), BaSyx does not provide support for resolving references to the referenced element. While this may not be a serious problem when following such links is not crucial, it is an obstacle for platform submodels such as services where we need to reference to a related service, resolve that and access the actual state. This absent functionality drove some of the structure decisions for our sub-models. Similarly, we use URIs to link sub-models and external AAS, in particular that the CLI/Management UI can provide information stated in the AAS.
- When deciding about the concept to realize, in particular AAS vs. sub-model, take the industrial production viewpoint where the AAS concept originates from and try to identify the asset that is described. If the modeling is about an asset (potentially provided by a different organization), typically an AAS is required. When detailing (own) information, often an own AAS is more adequate.

6 Platform Configuration Model

This section provides an overview on the IVML configuration model and the concepts used to model configuration options for the IIP-Ecosphere platform. We now give an overview of the configuration model, then, from a more pragmatic point of view, an insight into a simple example configuration as this is required for running the platform. Section 6.1 dives deeper into the configuration model by discussing applied modeling patterns, Section 6.2 provides more details on the structure of the configuration meta-model as well as on adjustments for a model managed by the platform, Section 6.3 discusses additional support for standardized protocols or connectors such as OPC UA, and Section 6.4 details the instantiation process. Section 6.5 discusses the container instantiation. We then turn to the development of applications based on the configuration model. Section 6.6 outlines some example applications shipped with the platform. Section 6.7 discusses the steps needed to create an application with the configuration and the instantiation process, Section 6.8 illustrates typical implementation project structures, and Section 6.9 illustrates default build sequences and their build commands. Finally, Section 6.10 summarizes service implementation considerations.

In essence, the configuration model mirrors the component hierarchy of IIP-Ecosphere and describes per component the configurable elements, their dependencies and constraints. IVML is the Integrated Variability Modeling Language [10] as realized by the EASy-Producer toolset [36]. The configuration model consists of three parts:

1. The **configuration meta-model** introducing the configurable elements, their structure, relations, properties and where adequate also consistency constraints.
2. A **platform configuration** based on the configuration model describing the configuration of a certain platform installation. Platform-specific structures (like services, service dependencies and service relations to form an application), but also the specific selection of alternative components, e.g., various transport protocols, service execution environments, container managers, are defined in the platform configuration. A platform configuration may introduce further, application/installation specific constraints.
3. A **valid platform configuration** complies with the configuration meta-model and fulfills all constraints. Such a valid platform configuration can be instantiated through an instantiation model, consisting of an instantiation process description (VIL, variability implementation language) and, where adequate, artifact instantiation templates (VTL, variability template language) [15]. In IIP-Ecosphere, both languages are used to instantiate a platform configuration into code and build specification artifacts, to execute and to package the created artifacts.
4. VIL and VTL can be used at **runtime to adapt the underlying system** [8]. These capabilities will be used in the last project year to allow for self-adaptation of the IIP-Ecosphere platform.

The configuration model is taken up by the configuration component (Section 3.11) and used for platform instantiation and runtime adaptation. The configuration component allows for high-level model operations.

As illustrated in Figure 55, the configuration meta-model reflects the layers and components of the IIP-Ecosphere platform, each given in terms of an IVML project. The most basic project (MetaConcepts) introduces even more abstract, i.e., meta-meta, concepts for generic adaptive software systems. These concepts are refined into IIP-Ecosphere specific concepts in the remaining models. The first IIP-Ecosphere specific model describes the DataTypes used in the platform, in particular PrimitiveType and RecordType consisting of files of DataTypes instances. Some specific primitive

types are defined in this model and frozen¹⁰⁸ already on that level. The remaining levels will be described as soon as they are realized.

The platform instantiation takes up the data types and turns them into language-specific artifacts, e.g., Java or Python classes. Similarly, corresponding serialization mechanisms to be used with the Transport component are generated. So far, there are no basic settings for the Connectors.

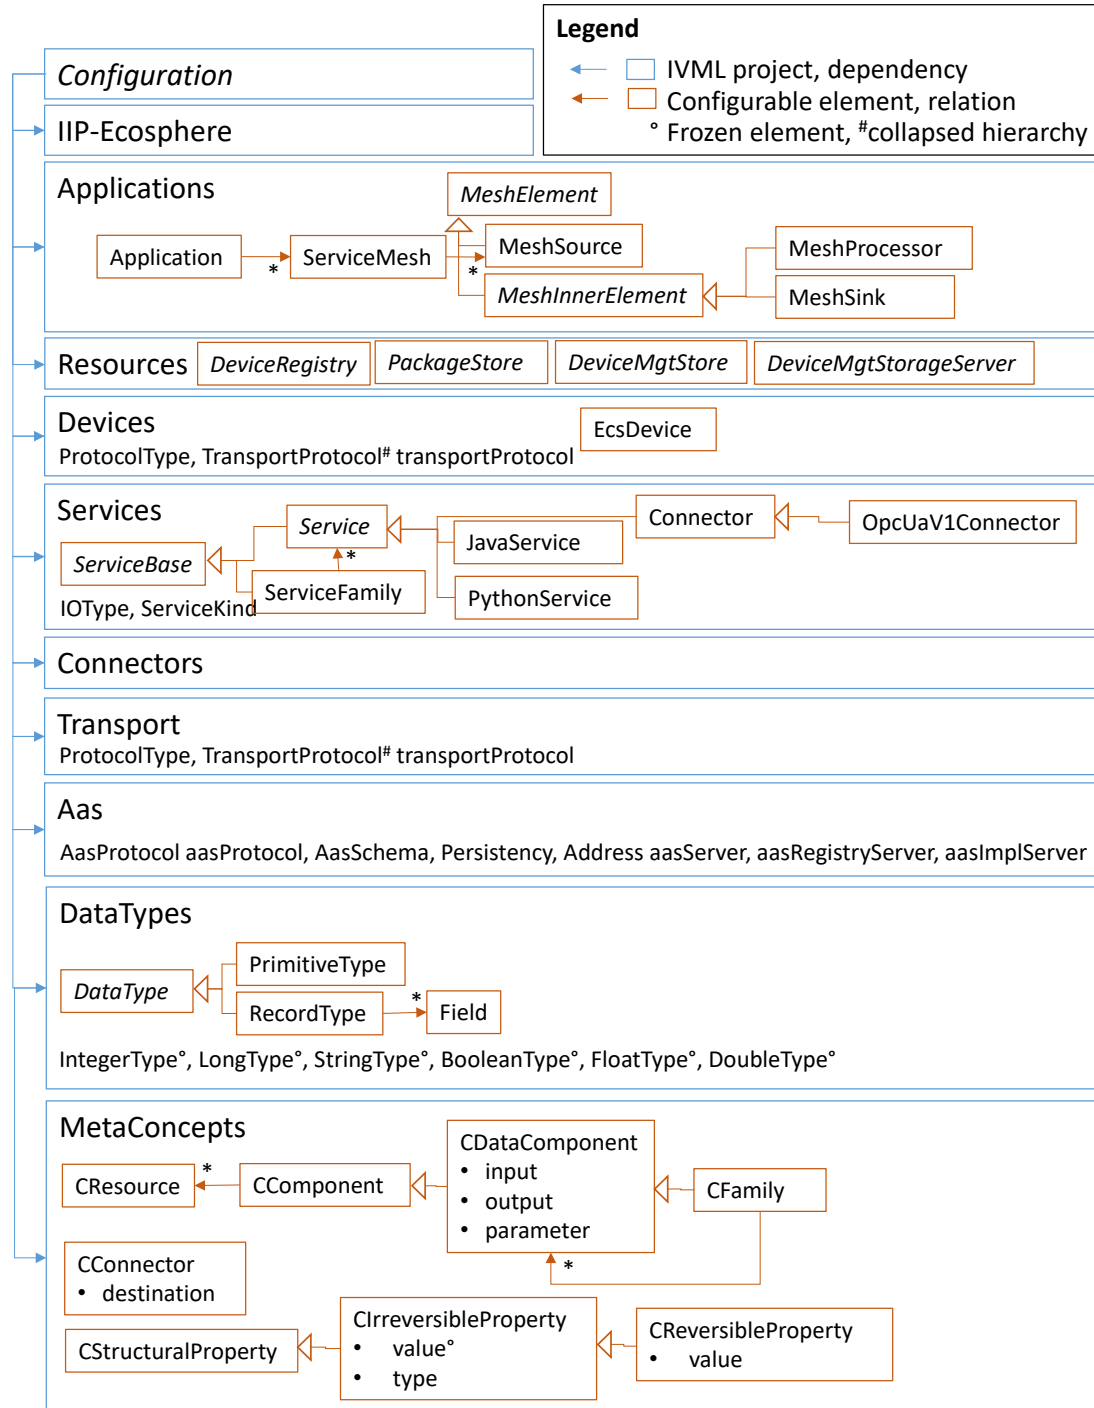


Figure 55: Illustrative structure of the IVML IIP-Ecosphere platform meta-model (simplified).

¹⁰⁸ Frozen elements cannot be modified outside the defining IVML project. Only frozen elements can be instantiated before runtime, while the remaining elements may be frozen later or remain changeable for runtime adaptation. The MetaConcepts model defines mechanisms to conditionally control the freezing and also the CReversibleProperty, which explicitly re-defines its value to remain unfrozen.

On the service level, several refinements of the IIP-Ecosphere service term are defined as configurable elements. The `ServiceBase` is abstract and contains information common to all services, e.g., name, id, version, description, input types, output types, service kind or operation mode (synchronous/asynchronous). Already the `ServiceBase` defines constraints prescribing which information must be present for which kind of service. Although we might use the service kind as hierarchy discriminator here, we opted for building a hierarchy along the implementation levels rather than the service kinds, as service kind differences can easily be handled by constraints while the implementation type is more important for the subsequent code generation. A `Service` is a refinement of `ServiceBase` and also the parent of language specific services like `JavaService` (e.g., detailed by a Java qualified class name denoting the implementation) or `PythonService`. A special kind of `Service` is a machine/platform `Connector`, representing the specific connectors implemented in the `Connector` component (only OPC-UA is shown here, similar elements exist for AAS, MQTTv3, MQTTv5 and AMQP). A `ServiceFamily` represents multiple, alternative but functionally equivalent services with the same input/output types. Service families steer the selection of alternative services at runtime. Although strange at a glance, a `ServiceFamily` (representing a concrete selection of one out of many services) is defined as a kind of service (it inherits from `ServiceBase`). This allows to transparently use a `ServiceFamily` wherever a `Service` can be used. From the configured services, the code generation derives implementation interfaces (Java, Python) and service stubs (Java) for the integration of non-Java service implementations.

The `Devices` module defines the properties of the ECS runtime, in particular the container manager to use. Moreover, it defines the `EcsDevice`, which represents an installed/connected device. In the next release we plan that `EcsDevice` instances steer the automated creation of Docker containers as well as the automated and optimized assignment of containers to resources.

The `Applications` module introduces one or multiple applications consisting of one or multiple `ServiceMesh` instances. A `ServiceMesh` is a directed graph (as introduced in Section 3.1.2) rooted by sources, linked by connectors/relations possibly leading to sinks. Each node in such a graph has an implementation in terms of a `ServiceBase`, which is refined to application-specific Java or services as well as platform-supplied services like connectors or pre-integrated services like the KODEX, the RapidMiner RTSA or the Trace2AAS service. Services declare their input and output data types, typically for forward or backward data flows (cf. Section 3.1.2). In the model, service properties are pulled up from service level to mesh level during model validation and allow for checking whether a flow graph is valid (through correctly sequenced input/output types of the services). During code generation, individual applications or alternatively all applications are processed, i.e., the service meshes are traversed and stream engine glue code for each node is generated. In the default case, Java classes with Spring Cloud Stream annotations are created and bound to the respective service interfaces. Based on the given implementation class names, the implementing services are dynamically instantiated, mapped into the respective AAS (via the `ServiceMapper` from the service environment) and made available for monitoring and management.

For building up AAS, e.g., to trace platform operations or to represent services, the configuration model also reflects basic vendor information that is required to instantiate respective AAS. This information can optionally be attached to an application or a service. Moreover, nameplate information items can be reused to increase consistency, e.g., if a “vendor” such as IIP-Ecosphere created multiple services.

Besides code artifacts also build specifications (Maven), assembly specifications, Spring application specifications, deployment descriptors, logging setting files, JSL specifications and, partially, test classes (for validating generated Yaml files) are created automatically. For the three major platform components, the platform AAS server (based on the `platform` component discussed in Section 3.13, currently without further services), the ECS runtime and the service manager, the basic AAS settings

as well as further settings are instantiated into respective Yaml application specification. Finally, the generated build specifications are executed so that for a complete instantiation, three platform artifacts and one combined Java/Python artifact per application is generated.

We do not provide a more detailed discussion of the concepts in the meta-model or the instantiation process at this point in time because both models are still in development and usually it is not expected that users of the platform modify the models. However, as long as there is no user interface, a user must be able to describe a platform configuration in order to perform an instantiation. Therefore, we briefly provide an insight into a simple testing model.

```
project SimpleMesh {

    import IIPecosphere;

    // binding annotation omitted

    // ----- component setup -----

    serializer = Serializer::Json;
    // serviceManager, containerManager are already defined

    aasServer = {
        schema = AasSchema::HTTP,
        port = 9001,
        host = "127.0.0.1"
    };

    // ...

    // ----- data types -----

    RecordType rec1 = {
        name = "Rec1",
        fields = {
            Field {
                name = "intField",
                type = refBy(IntegerType)
            }, Field {
                name = "stringField",
                type = refBy(StringType)
            }
        }
    };

    // ...

    // ----- individual, reusable services -----

    Service mySourceService = JavaService {
        id = "SimpleSource",
        name = "Simple Data Source",
        description = "",
        ver = "0.1.0",
        deployable = true,
        asynchronous = true,
        class =
            "de.iip_ecosphere.platform.test.apps.serviceImpl.SimpleSourceImpl",
        artifact = "de.iip_ecosphere.platform:apps.ServiceImpl:" + iipVer,
        kind = ServiceKind::SOURCE_SERVICE,
        output = {{type=refBy(rec1)}}
    };
}
```

Figure 56: First part of a simple platform configuration.

Figure 56 depicts the first part of a simple platform configuration used for testing. A model is defined in terms of IVML, a textual DSL for variability modeling. Each model is surrounded by a project namespace, here named `SimpleMesh`. Within that namespace, first model imports are stated, here an import of the IIP-Ecosphere configuration meta-model (`IIPEcosphere`). After this header, the first configuration value definitions are stated, typically as value assignments to typed variables (a typed variable indicates a configuration option in IVML). Typed variables can form complex types that we call compounds in IVML. Here, the serializer is defined to be `Json`, an enumeration literal for serializers defined in the meta-model. Then the global `aasServer` receives its schema, port number and host name (similarly but not shown for AAS registry and local AAS implementation server). Next, we define the application datatypes, typically records.

While the variables discussed before are pre-defined by the meta-model, the data type is now given in terms of an own variable named `rec1` of type `RecordType` (defined in the meta-model as a compound, not illustrated here). A record has a name (turned e.g., into a Java class name during instantiation) and field, each with a name and a type. Types are references (stated by `refBy`), i.e., we define a link to an already defined variable, here the pre-defined `Integer` and `String` type.

Following the definition of the variable `rec1`, we then introduce a Java service, a hand-crafted data source (for testing, it will create arbitrary data of type `rec1`). The source is described by its identification, its name, an empty description, a version, whether it is deployable, whether it is a synchronous or asynchronous service and its implementation class located in the given Maven artifact. Please note that we use here the implementation version of the platform defined by the meta-model in the variable `iipVer`. The service is a source service (one of the four main service kinds) and its output is constituted by one record, namely `rec1`. In fact, multiple types can be given, all in terms of a structured type currently just having a type field (to be extended later), therefore the double brackets, the outer one for a collection instance, the inner one for the structure type.

```
// ----- application and service nets -----

Application myApp = {
  id = "SimpleMeshApp",
  name = "Simple Mesh Testing App",
  ver = "0.1.0",
  description = "",
  services = {refBy(myMesh)}
};

ServiceMesh myMesh = {
  description = "initial service net",
  sources = {refBy(mySource)}
};

MeshSource mySource = {
  impl = refBy(mySourceService),
  next = {refBy(myConnMySourceMyReceiver)}
};

MeshConnector myConnMySourceMyReceiver = {
  name = "Source->Receiver",
  next = refBy(myReceiver)
};

MeshSink myReceiver = {
  impl = refBy(myReceiverService)
};
```

Figure 57: Second part of the simple platform configuration.

The second part of the example in Figure 57 defines an application with a simple service mesh. First an application is defined, again with identification, name, version and empty description. Then the service meshes are stated, here a single reference to `myMesh`. `myMesh` potentially consists of multiple sources, we just have `mySource` as source mesh element. `mySource` uses the previously defined `mySourceService` as implementation, as well as the next mesh element in terms of a mesh connector/relation. A synchronous source may also define a polling interval. Currently, mesh connectors have just a name but further properties may follow (otherwise we could directly reference mesh elements among each other). The mesh connector links further to the receiver, which states its implementation as `myReceiverService` (similar to `mySourceService` but not shown here).

```

    freeze {
        aasServer;
        serializer;
        // ...
        .;
    };
}

```

Figure 58: Final part of the simple platform configuration.

The final part is important for the instantiation. For various reasons, variable values defined in IVML are not per se considered final, rather they can be overwritten in importing modules/project. Turning such a configuration into code is problematic, in particular if code parts are deleted based on non-final decision (deleted parts are usually deleted). Thus, IVML has the notion of freezing variables. Frozen variables are considered final and can be instantiated safely. Figure 58 illustrates the freezing of this model. Within the freeze block, first variables from the meta-model that have been configured are listed for freezing. Finally, every variable declared in this project (shortcut “.” like in a command shell) is frozen. Typically, in systems with dynamic instantiation at runtime, freezing is conditional, i.e., stated variables are filtered according to a given condition. In the original model used for testing, this condition is based on the so-called binding time, the latest time when a decision must be made (here compile time). As we just aimed at explaining how a platform configuration looks like, we intentionally left out the required attachment of binding times at the beginning of the model and the freeze condition here. Ultimately, Figure 58 ends with the closing bracket for the namespace of the `SimpleMesh` project.

Although the configuration shown here looks pretty structural and might be represented in any other nested configuration language, we did not detail the validation constraints that are imposed by the meta-model, e.g., that services are configured correctly and services meshes fit together. For now, the constraint setup is initial and several constraints are currently missing. However, the already defined constraints can quickly lead to validation errors issued by the EASy-Producer reasoner. This validation is important, as an invalid model typically leads to invalid artifacts that, e.g., cannot be compiled. Work is still needed here to make the validation messages more domain-specific and user friendly.

In summary, the code generation based on the IIP-Ecosphere configuration model creates more than 14 different types of artifacts (Maven XML, assembly XML, Java source, Python source, application Yaml, logging XML, Java test code, Windows batch/Linux shell startup scripts, Linux/system service descriptors, README files, Broker setup specifications, Docker files, Type script files, Angular environment setups), which leads to different types of artifact structures, e.g., various forms of Java code. The number of generated artifacts varies with the number of services/mesh elements defined per application/platform configuration.

Besides a brief explanation of a configuration model, it is probably relevant to the reader to have executable examples or tutorials at hands. We will cover this topic in Section 6.5.

6.1 Modeling Patterns

As shown in Figure 55, the platform configuration model consists of several layers reflecting the architectural layers of the platform. Each configuration model layer defines the decisions to be made, typically either a) using basic IVML types b) refined compound types if the alternatives have detailing properties or c) a more detailed structure of own types to model service and app decisions. This section dives a bit deeper into the IVML platform configuration model.

```
decision = X::Alternative2;

enum X {Alternative1, Alternative2, ...};

X decision = X::Alternative1;
```

Figure 59: IVML model pattern for simple alternatives without detailing properties.

Figure 59 shows the IVML model pattern to represent simple alternatives that do not need to be detailed further, e.g., the transport layer serializer format. The lower box in Figure 59 illustrates the model layer, the upper box the specific platform configuration. The alternatives are modeled (in the lower box) as the enumeration type *X* listing all potential alternatives. The declaring model layer also defines a variable representing the respective *decision* and assigns a default value to ease creating a configuration. The configuration (upper box) overwrites the value to indicate that a different alternative shall actually be included into the platform instance. It is important to note that this pattern does not allow for openness as IVML enums are fixed and cannot be extended later, e.g. in importing IVML modules.

```
decision = Alternative2 {
    // assign properties as needed
};

abstract compound X {
    Type p = default;
}

compound Alternative1 refines X {
    //optional further properties, constraints
}

// further alternatives, constraints

X decision = Alternative1 {
    p = default1
    // assign further properties
};
```

Figure 60: IVML model pattern for alternatives with detailing properties.

Many alternatives demand further information when selected, e.g., the transport protocol, the S3 storage client/server or the AAS client/server settings. In this case, we model alternatives as IVML compounds, an abstract base compound defining a common type for all alternatives and refining compounds for the individual alternatives. The base type, in Figure 60 the compound *X*, defines properties that are common for all alternatives, e.g., a server port, usually with default values, while the individual alternatives such as *Alternative1* may add further properties. Each type representing a specific alternative can define constraints that become active only if that specific alternative/type is used. The alternatives may override the default values by re-declaring the properties with the same

name/type. As in the first pattern, the declaring model layer defines a variable representing the respective decision, assigns a default instance of *Alternative1* including specific values for the properties. The configuration in the upper box may then assign a more specific value, here *Alternative1* and properties. Please note that property values can be derived from, e.g., common global variables to increase consistency. Moreover, modeling alternatives via compound refinements allows for openness, as further refining alternatives can be defined on any upstream model level, i.e., this IVML model pattern is appropriate for alternative components and plugins contributed by the user. This form of openness must be adequately taken into account in the instantiation process.

In several situations, the configuration model must remain open for extensions by the user or by third parties (supporting the idea of a service store in [41]). Typical extensions are service or application definitions as well as device types (acting as templates for device-specific container instantiations). Akin to the configuration model, also the instantiation process must be extensible, e.g., to perform service-specific generations when an externally provided generic service shall be integrated. The EASy-Producer languages IVML, VIL and VTL offer mechanisms to support such requirements. One basic mechanism is dynamic dispatch operations, for user-defined constraints as well as for generation functions. In dynamic dispatch, the actual operation to be executed is dynamically determined based on the actual types of all parameters (an extension of first-parameter polymorphic execution in object-oriented programming languages like Java) [14]. Basically, dynamic dispatch allows to consider later extensions of the actual model by refinement and, thus, provides a basic form of openness to refined, but yet unknown types. A second mechanism allows for a dynamic model structure through wildcard imports (similar to wildcard imports in Java), i.e., model parts that are not known at modeling time can dynamically be added to the model structure. However, imports just make model elements known to other model elements and do not influence the dynamic dispatch mechanism. Thus, the EASy-Producer languages provide a special import statement, that allows to extend the dynamic dispatch in the model at hand by dynamically loaded models. In essence, externally provided model parts are dynamically loaded into the model structure and can hook themselves into the constraint and generation operation through specialized types defined in these models.




Figure 61: Model structure for openness and extensibility.

For the aforementioned situations, this requires a specific model import structure as shown in Figure 61. Let *X* be the part of the model for services, applications or devices, e.g., *X* could be “Devices”. To enable the kind of openness as explained above, one model module must define the basic types and operations the dynamic extensions shall hook into (module *XBase*, for the devices example the name would be *DevicesBase*). The extensions are modules that import the base module and add own types and operations. Their names are composed from the respective model part, i.e., *X*, the infix “Part” and the individual name of the extension, e.g., *DevicePartPhoenixContact*¹⁰⁹ for specific device types representing the AXC PLC/edge series of Phoenix Contact. Finally, some modules must use the extensions. Such a module must be an extension of the *XBase* and dynamically import the *XPartY* in a way that they hook into the dynamic dispatch of *X* (and transitively extend *XBase*). In the IIP-Ecosphere model, this is the module *X*, in the example called “Devices”, which usually is empty except for imports and the model extension statement. This module is then imported by further modules of

¹⁰⁹ Depending on the naming, the prefix may be adjusted deliberately for a more “speaking” name.

the configuration model and implicitly introduces the dynamic extensions. At runtime, further extensions can be added, but currently only by re-loading the entire model.

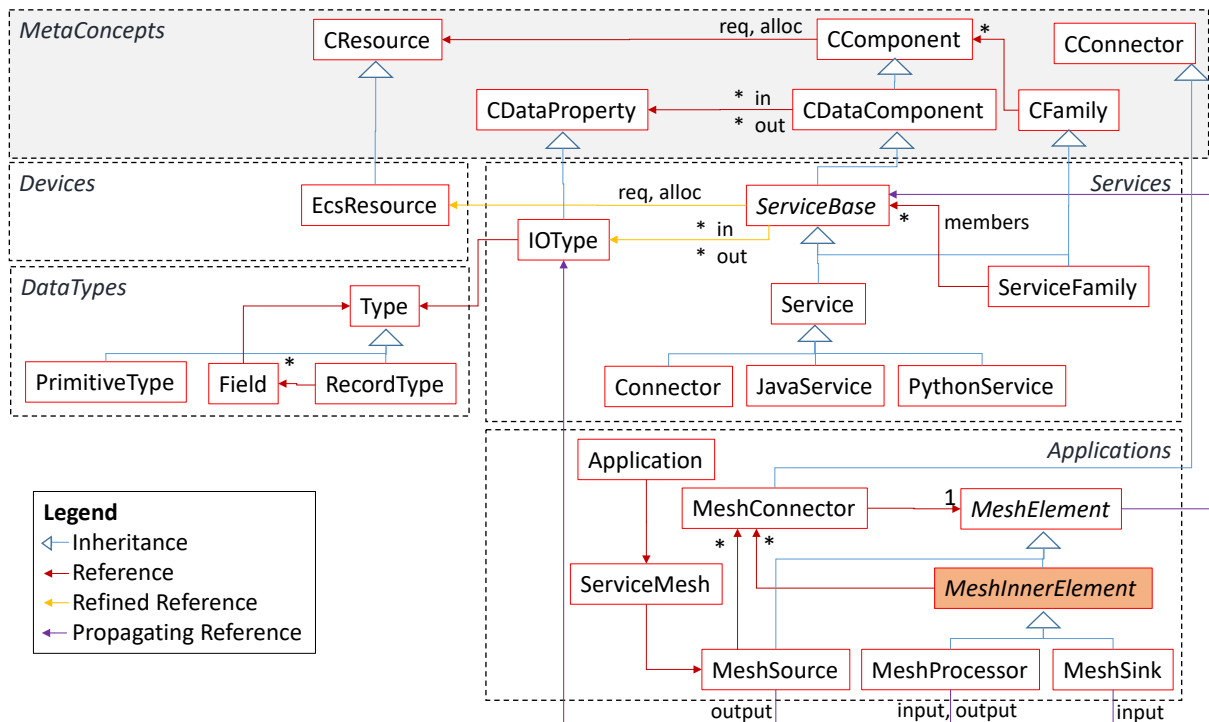


Figure 62: Meta-model concepts for defining services and alternatives.

The configuration of user-defined applications that are executed on top of the platform, the involved services, their data paths and the resources to execute the services on requires more information and, thus, is more complex than the two IVML model patterns discussed before. The most relevant configuration concepts for applications are illustrated in terms of the UML-like class diagram in Figure 62. We target the following aims:

1. Configure **re-usable services** (the *Services* module in Figure 62) and their properties, potentially families of semantically equivalent services that can be exchanged at runtime, e.g., alternative AI services. Services include those provided and integrated into the platform as well as user-supplied services.
2. Represent **data transformation and mappings** to reduce the effort of manual coding in standard situations (in the sense of “low code”, R19f). Currently, we apply such data transformations in particular to integrate (machine) connectors. Figure 63 illustrates the applied data mapping approach. For a connector we specify a machine- and a platform-side I/O data format, usually a record of named/typed fields. Fields with same (nested) field names are mapped onto each other in both directions, machine-to-platform and platform-to-machine. Fields that cannot be mapped are either ignored, i.e., projected out, or left uninitialized. To fill individual fields, assignment expressions between both sides can be stated, allowing for simple data transformations, e.g., unit calculations. For a model connector (cf. Section 3.6.3), the given data formats can be turned by the generator into model paths and data can be obtained and transformed automatically. For a channel connector (cf. Section 3.6.3), the input is always binary. Here, the input parsers from the Connectors component can be specified in the model to turn the data into named/indexed fields that are mapped by default to the machine formats and, thus, can further be used as for the model connectors. In the opposite direction, output formatters can be applied. If no platform-supplied parsers/formatters fit to the data at hands, own Java components supplied as Maven

component can be specified. Similarly, the entire mapping process can be bypassed by own serializers or model adapters.




Figure 63: Overview of low-code data mapping for connectors.

3. Configure **physical and logical compute resources** the services are executed on (the *Devices* module in Figure 62), although it is important to emphasize that configuration instances of such resources shall be created and reflected into the configuration by the device management at runtime.
4. **Compose connectors and services to applications** (the *Applications* module in Figure 62) so that one service can occur in multiple applications and that the data paths within an application are defined and can be instantiated automatically.

In more details, the most basic module in Figure 62 is *MetaConcepts*. Although concepts and properties defined in this module could also be introduced in the dependend modules shown in Figure 62, the aim of *MetaConcepts* is to represent generic concepts of configurable runtime-adaptable systems. Thus, *MetaConcepts* introduces basic notions of resources (CResource), components (CComponent), families of components (CFamily) and connectors among components (CConnector). As these concepts define properties using these types (and the connectors even of the top-most IVML type Any, therefore no associations in Figure 62), which must be re-defined in upstream modules, e.g., in the *Services* module.

From the generic *MetaConcepts* perspective, we now turn to the IIP-Ecosphere specific configuration concepts. The *Services* module currently just introduces the notion of a device, where additional properties will be added in the future. The *DataTypes* module introduces the ability to express types that are reflected in current programming languages, such as “primitive” types like String or Integer, but also more complex, composed types (called RecordType). These types are used in the *Services* module to specify the inputs and outputs of individual services. In this module, specific service types (Connector for the connectors in Section 3.6.3, JavaService for services implemented in Java and PythonService for services implemented in Python) are defined and also service implementations shipped with the platform (AAS, OPC UA and MQTT connectors from Section 3.6.3) are defined as configuration instances. Akin to *MetaConcepts* a refined type for user-defined service families that can act on behalf of a fixed, individual service is introduced.

On top of these configuration layers, the *Applications* module defines graphs of services, called service meshes. An Application consists of one or multiple ServiceMesh instances, and, in turn, a service mesh starts at one or multiple sources (of type MeshSource). Sources are linked via MeshConnector instances to processor or, ultimately, sink nodes. In contrast to the IVML model used in the FP7 QualiMaster project [14], we cannot restrict inner nodes to processors and sinks, as processors may have backward flows to control machines via connectors.

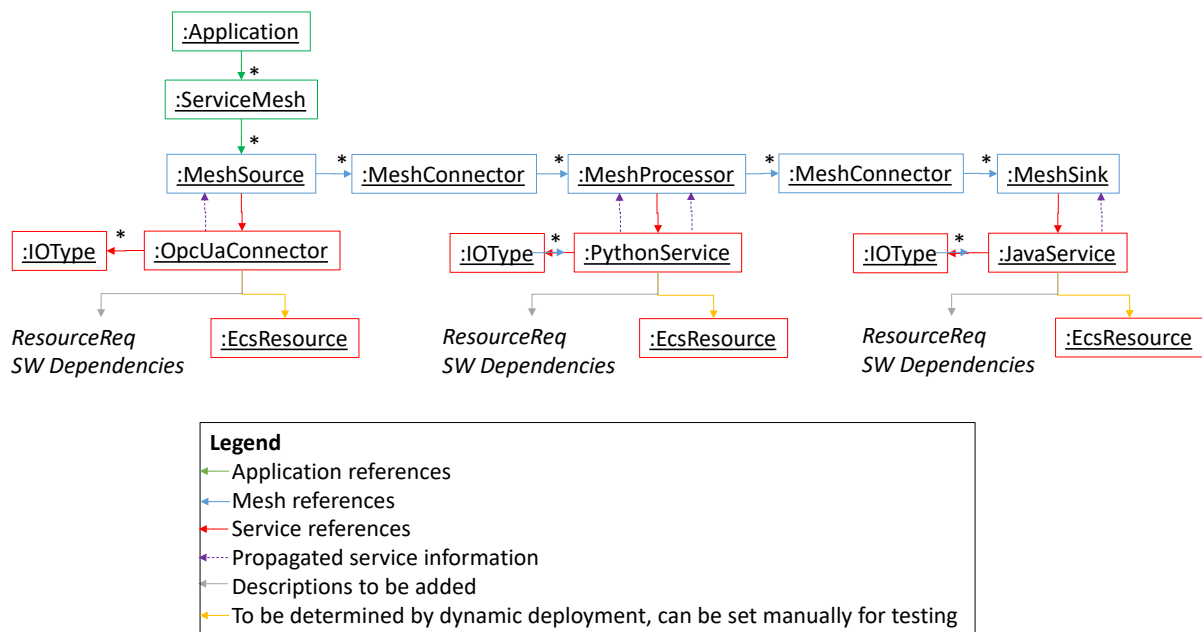


Figure 64: Instance view on an IIP-Ecosphere platform application.

As illustration of these concepts, Figure 64 shows how instances of the aforementioned types can be linked together (backward flows are not shown in Figure 64). The Application consists of one ServiceMesh, which, in turn, consists of a chain of three services, a source, a processor and a sink, all linked by instances of MeshConnector. The source is implemented by an OPC UA connector, the processor by some Python implementation, e.g., an AI algorithm, the sink by some Java implementation, e.g., a database. Each of these services has its own input/output types, which must comply with the predecessor/successor services along the graph constituted by the service mesh. Further, each service is (at latest at startup time of the application) deployed to a certain resource, e.g., an edge device. To determine adequate resource candidates, currently descriptions of resource requirements and software dependencies to be installed into hosting containers are developed.

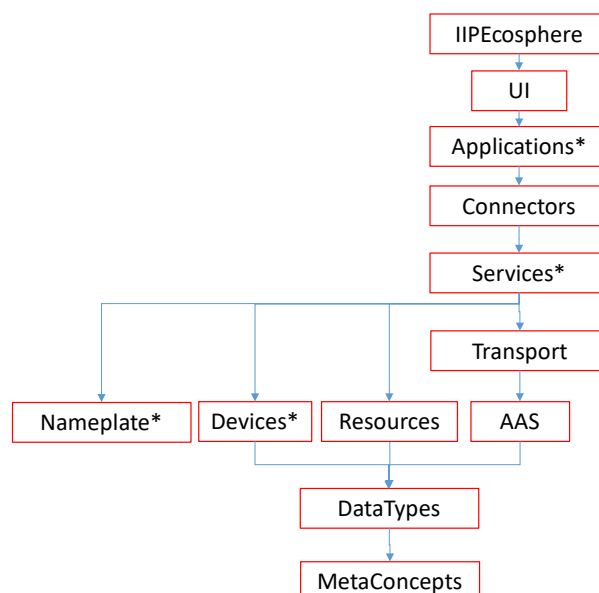


Figure 65: Structure of the IVML configuration meta-model for IIP-Ecosphere (transitive dependencies omitted).

6.2 Configuration Model Structure

The structure of the configuration meta-model (Figure 65) follows the overall structure of the IIP-Ecosphere platform architecture with some additions and a different layering/dependency approach

due to pragmatic reasons. The individual modules define the configurable properties and types for the respective architecture layer, component or aspect. The bottommost module `MetaConcepts` stems from an attempt to capture the basics of an adaptive software system and is included here for evaluation purposes. The `DataType` module defines primitive and extensible data types used for specifying input/output types of services and connectors. On the next level, AAS server/implementation properties, AAS Nameplates with default instances for IIP-Ecosphere, Devices including capabilities and requirements, and Resources (including resource management and monitoring) are defined. Then, the actual transport protocol and its setup (including authentication and transport layer security) are specified. Services (including Java, Python services), reuse all these configuration types and form a basis for the more specific Connectors. Applications and their service meshes are defined on top, imported by UI (settings) and finally the top-level `IIP-Ecosphere` module for global technical installation settings such as installation paths. Module names suffixed with a `*` are extensible through the import mechanism explained in Section 6.1.

An actual copy of the meta-model is included in each implementation/example project after downloading the model through Maven. In these projects, the actual configuration is typically defined in one or two modules to focus on the specific properties of the example and to ease gaining an overview.

When the platform takes over control of its own model, fixed structures are needed so that the configuration component can store the configured settings, service instances and application meshes. This form of models is currently in testing and will be enabled as soon as the management user interface provides respective editor functionality, in particular a graph-based service mesh editor based on the information provided by the configuration AAS. The model structure required then is depicted in Figure 64.

In this setup, the topmost module is the `PlatformConfiguration` storing settings that override global non-frozen configuration options. Service instances are stored in `AllServices`, related type definitions for input/output specifications in `AllTypes`. In turn, `AllTypes` relies on `AllConstants`, containing e.g. server host names or commonly used port numbers. The application instances (`ApplicationPartX`) are stored in individual extensions pointing to their linked service mesh parts (`ServiceMeshPartX`). The top-level meta-model module `IIP-Ecosphere` and transitively imported modules are linked through `AllConstants` into the managed configuration model. Further, `TechnicalSetup` defines the fundamental technical abilities of the platform, e.g., transport protocol, monitoring or management UI setup. In turn, `TechnicalSetup` may rely on the constants in `AllConstants`.

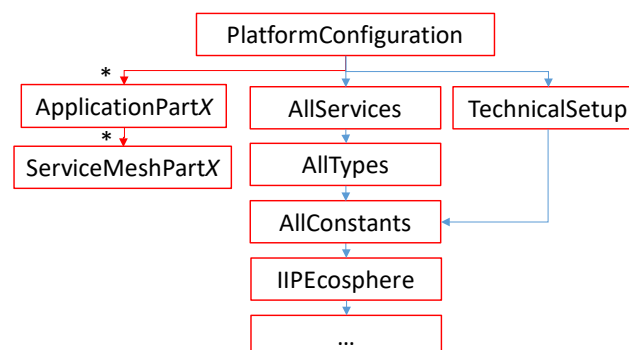


Figure 66: Managed model structure.

6.3 Support for Standardized Connectors/Protocols

Defining the input/output data types for complex, nested data structures can be a complex process. For standardized protocols/information models such as OPC UA or AAS, more and more data structures

with standardized fields and semantics are created. For a standard-based platform like the IIP-Ecosphere platform it is important to take up such approaches and to ease the use of standardized data structures, thus, supporting the platform user in creating applications.

As an example, we turn now to the OPC UA Companion Specifications, a set of standardized models for OPC UA. Currently, more than 50 such models have been specified and further models are in preparation. Over time, also the defined model structures are evolving. Thus, manually translating OPC UA Companion Specifications could be an initial approach, which will not turn out to be sustainable. Fortunately, the OPC Companion Specifications are available in a machine-readable XML format, which can automatically be translated into IVML [10] and then used further in application configurations. We demonstrate this approach by an automated model translator [5].

The model translator reads an OPC Companion Specification XML file and translates it into IVML using a base meta-model, which extends the IIP-Ecosphere configuration meta-model, in particular the DataTypes module. The created IVML files, one per companion spec can be imported into an own application model when needed. Besides the main types representing a Companion Specification, also declared subtypes can be used in custom applications, in particular as input/output to generated OPC connectors.

The approach was successfully validated with all 55 available OPC UA companion specs. Valid IVML models were produced for all companion specs and valid connector code was generated. The largest specs, e.g., the Tobacco Machine Communication (TMC)¹¹⁰ or the IEC61850-7-4 spec triggered a restructuring of the generated connector code to cope with the more than 120 KLOC XML specification. As far as available in the VDW UMATI OPC test server, 3 of the generated connectors were successfully functionally validated against an OPC reference implementation [5].

6.4 Platform Instantiation Process

After successfully configuring a platform and the apps to run on the platform, the configuration must be instantiated. This happens through further languages of EASy-Producer [15], namely the Variability Instantiation Language (VIL) to express the control over the instantiation process and the Variability Template Language (VTL) to modify or create artifacts of a certain type, e.g., XML or Java code files.

Figure 67 illustrates the steps that are executed during the instantiation. The IIP-Ecosphere VIL model defined three major entry points, which are available through the PlatformInstantiator tool, namely

- `generateInterfaces` generates the interfaces of the declared applications as a basis for the realization of user-defined (non-platform provided) services as well as default basic implementations of the service interfaces, e.g., to ease the implementation of parameters and data ingestors.
- `generateAppsNoDeps` instantiates the applications but intentionally leaves out all implementation dependencies. In addition to `generateInterfaces`, this provides also access to implementing base classes. However, the user-defined service implementations (whether they already exist or not) are not considered during the instantiation process.
- `generateApps` for the instantiation of (currently) all defined apps for a platform including all required dependencies. The resulting service artifacts are packaged to be executable, although of course bugs and errors may be located in the application logic.
- `generateBroker` generates an example broker/service instance the configured transport protocol. Typically, we rely on the broker implementations that we use during regression

¹¹⁰ <https://reference.opcfoundation.org/TMC/v200/docs/8.1>

testing. These instances may not be intended for production code, but they are helpful for the first setup or for examples.

- `generatePlatform` for the instantiation of the platform components
- `main` (not shown in Figure 67) which executes all aforementioned entry points, in particular for testing.

One could generate the interfaces for applications in one step with the application code and packaged as part of the same artifact, e.g., using different Maven classifiers. Although this approach is easier to realize, it hinders service reuse and may cause cyclic builds as it seems that Maven does not separate between dependencies of the main and classified artifacts. In contrast, we rely on a separation into an artifact that contains the generated interfaces for all applications defined for a platform and, in individual Maven artifacts, the service implementations. It is important to mention that the service implementations can be organized along applications, but, to facilitate reuse, an artifact can also realize individual services, in extreme a single service. Figure 67 mainly illustrates the separated instantiation approach.

For the instantiation of the **application interfaces**, we first iterate over all data types declared in a platform configuration and create their Java and Python realization (`JavaType`, `PythonType`). Moreover, we create the related serializers based on the declared types, for Java in order to realize the platform transport wire format (here JSON is just one alternative) and for Python a JSON-String Serializer to link a Service into the Python Service Environment. For all services in the platform configuration, we generate the service interfaces (`JavaServiceInterface`, `PythonServiceInterface`) and where feasible a basic implementation for service parameter and ingestor handling (`JavaServiceBaseImpl`)¹¹¹. Please note that, as discussed in Section 3.7.3, Java and Python Service Environments are similar, but also differ in the required level of programming, which is reflected in the different instantiation steps. Thus, Java services require a more complex code generation due to the direct integration into the service execution engine than Python services that are “just” executed by the Python Service Environment, which, in turn, is integrated through Java code into the service execution. At the end of the creation of the application interfaces, we create a Maven assembly descriptor for the Python interfaces, a Maven build specification that creates the deployable artifacts as well as an Ant script to execute the deployment in the continuous integration.

For obtaining the applications (integrating the handcrafted service code), we iterate over all application specifications and all their declared service meshes. Here, we create the code to bind the respective service into Spring Cloud Stream (`JavaSpringCloudStreamMeshElement`) and for further integrations a stub class based on the service interface to transparently integrate non-Java implementations into the mesh. Depending on the actual service, also further artifacts and assembly descriptors may be generated, e.g., as shown in Figure 67 for KIPROTECT KODEX. Furthermore, for each app, the deployment descriptors (if specified in the configuration), a starter class (for registering the mesh elements to the service framework and for registering the serializers), and the Maven POM file (`AppMvn`) are created. The POM file is executed, ultimately creating several artifacts, including one containing the programming interfaces as well as one representing the application-specific service artifact (including dependencies, service code etc.). For deploying the artifacts to Maven in our continuous integration environment, also an ANT file is created.

¹¹¹ To allow for service implementations that are not based on the basic implementation, the default values of the service parameters are set after the instantiation of the service through the reconfiguration operation, i.e., the parameter values are not available during the execution of a constructors but shortly after, usually when the first data arrives for processing.

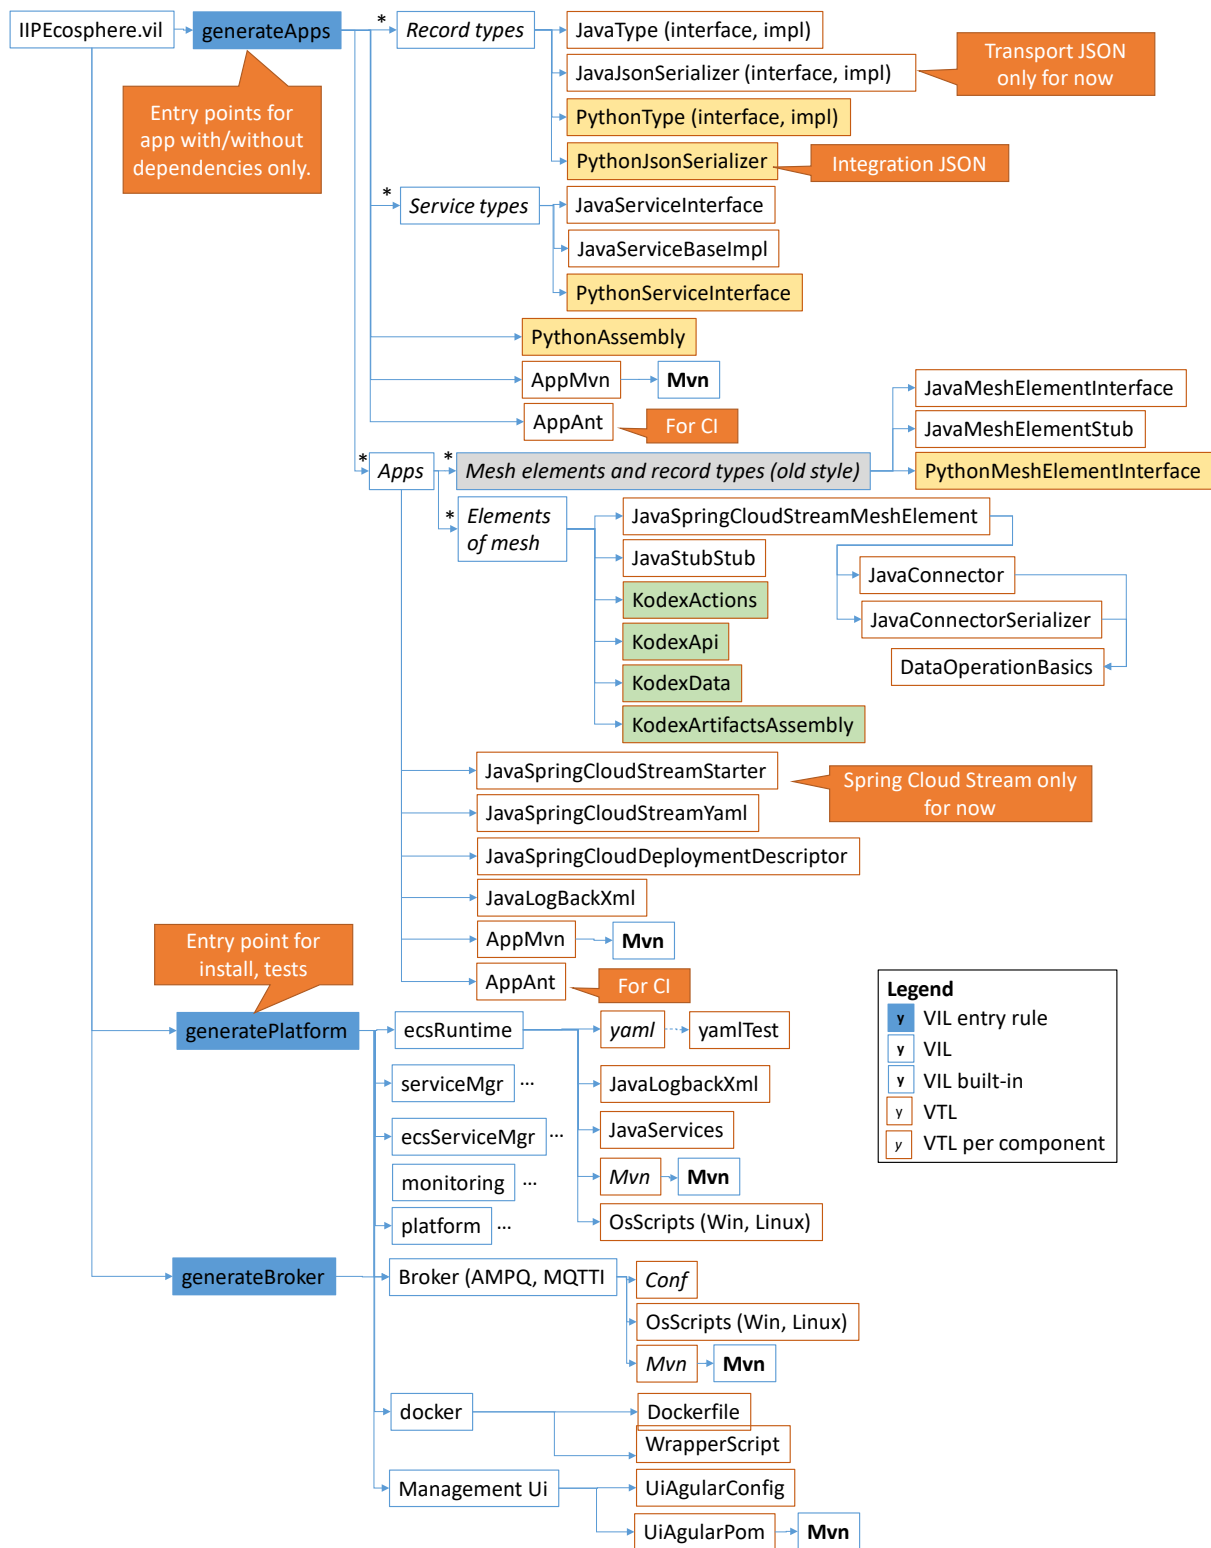


Figure 67: Overview of the platform instantiation process.

On platform level, the instantiation process creates packaged artifacts containing the ECS-Runtime, the service manager, a combined version of ECS-Runtime and service manager as well as the central platform services. For each of these components, first the application setup file (`yaml`) and a test class to validate the Yaml file are created. Then the logging configuration, the selected JSL service descriptors and ultimately a Maven POM with the respective components selected in the platform configuration are created. The Maven POM is executed, ultimately creating the respective artifact, more precisely, folders containing all required dependencies. Moreover, for starting the components,

operating-system specific scripts for Windows and Linux (here also descriptors for automatically starting the components as operating system services) are created. Similarly, the configured transport protocol leads to the instantiation of a corresponding (test) broker, which is also an important prerequisite to generically run examples from an IVML configuration. Further, Docker containers are instantiated based on service, dependency and installation information as well as the download/instantiation of the platform management UI is performed.

The code and artifact instantiation for Spring Cloud Stream currently applies a special strategy for connecting services. This strategy stems from practical experiences, where asynchronous data connections between services did not (always) work out as expected. Already reported Spring Cloud Stream problems may be the root cause here, but we were not able to identify the problem. Thus, we currently route almost all asynchronous connections directly via the transport layer. If we would fully rely on Spring Cloud Stream, a binder plugin would transparently do this for us, which failed in some cases. One particular exception are connections from asynchronous services to synchronous services, which cannot be realized without Spring Cloud Stream mechanisms. We plan to revert to Spring Cloud Stream based on more regression and use case tests in the future.

Specific services such as KODEX, RTSA or the upcoming Flower-based federated learning integration are extensions to the configuration model and the code generation and are represented in their own modules that are dynamically loaded into the platform configuration/code generation.

6.5 Container Instantiation

Virtualization of platform or application functionality intertwines technical requirements with convenience/ease-of-use. On the one side, several kinds of devices, in particular edge devices, do not allow for extensive software installation outside their own ecosystem. In such environments, Java/Python in general or required versions of these languages or their libraries in particular may not be available and an installation may not be permissible. Due to the trend of virtualization in IIoT, in particular for edge devices, but also due to the availability of Docker and reasonable CPU and memory capacities (sometimes even GPU or TPU processors), virtualization becomes a mandatory functionality for the IIP-Ecosphere platform (see also [13]). On the other side, complex services require non-trivial dependencies, which, in particular for Python, typically requires a physical installation of the required packages, which, in turn, depending on the underlying operating system or Linux distribution, may demand the installation of native libraries, the execution of operating system installation procedures or C code compilers. Such operations may not be permitted on a target system, but, more importantly, separating specific dependencies among different (versions) of applications or uninstalling/cleaning up such dependencies without virtualization capabilities may end in a nightmare.

Creating adequate containers for an IIoT application requires technical knowledge and, dependent on the actual (performance) goals, may involve non-trivial tradeoff decisions, e.g., pre-installed vs. dynamically installed dependencies or regarding the layering of the target containers, in other words, what to do first while creating the container in order to maximize with respect to desired system properties (transfer time, upgradability, adaptability). For this purpose, the IIP-Ecosphere platform aims at automatically creating the required containers based on a set of basic container creation strategies. As ideally no human is involved in this process, this allows for application and device specific containers (e.g., vendor builds of certain libraries) instead of one-size-fits-it-all containers as often created in other/industrial approaches, typically for virtualizations in cloud settings where certain resources do not matter so much [11, 12].

Container instantiation is the process of creating container images that include an application with all the dependencies that it needs to run an application at hands. We primarily focus on the Docker

container technology (mentioned but not required in [13]), but allow for extension by other approaches, e.g., LXC. Container instantiation happens during the platform instantiation.

The instantiation process is enabled when the configuration model for an application sets the variable `createContainer` to `true`. Then, based on the technical setup of `ContainerType` in the configuration model, three different types of container images are created. Successfully created container images are stored in the respective container registry as defined in the configuration model. The registry could be a public registry such as DockerHub, but in particular also private, locally installed registries are supported in order to allow for applications that are based on licensed or IPR protected components.

Currently, as discussed in more details in [40], the platform supports the automated creation of three types of container images instantiation (based on the configured value for `ContainerType`). The container types are illustrated in Figure 68. These types are

1. `Ecs_Svc_App` to create a container that contains an ECS runtime and a service manager running as separate processes. The container contains the dependencies of all services for the configured application(s) as well as a local communication broker to facilitate local transport-layer data communication among application services.
2. `EcsSvc_App` to create a container that will have an ECS-runtime and a service manager running as one single process. Although it is advisable to run these two central services as own processes to increase resilience, a single process allows for saving memory resources. As for `Ecs_Svc_App`, the container includes the dependencies of all services for the configured application(s) as well as a local communication broker to facilitate local transport-layer data communication among application services.
3. `C1Ecs_C2Svc_App` to create two separate containers: one container will have an ECS-runtime and the second container will have a service manager, the application dependencies and a local broker for local transport-level communication among services.

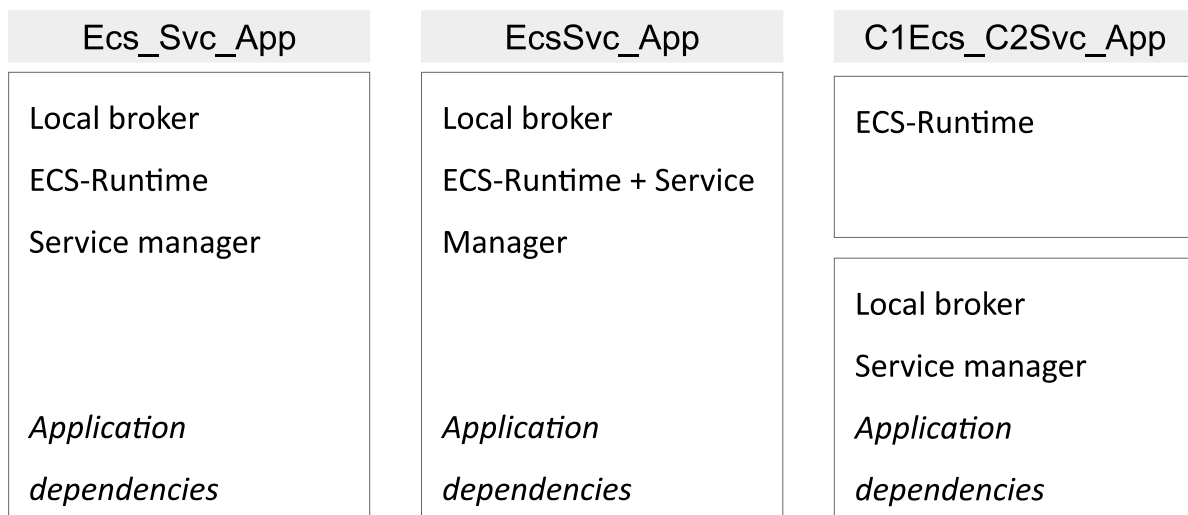


Figure 68: Container types and contained services/parts as supported by the IIP-Ecosphere platform.

After creating the images and pushing them to the registry, a descriptor for each Docker image is created specifying properties of the container image. Besides descriptive characteristics of the container, the descriptor in particular contains the container startup options (environment variables, exposed network ports, utilized volumes, Docker-out-of-Docker settings for the `Ecs` or the `C1Ecs_C2Svc_App` container type) and the location from where to obtain the container, e.g., an image name pointing to the local container registry. These descriptors are stored in the artifacts folder

of the platform as specified in the configuration model and can be obtained by the management user interface (to display the available containers), but in particular by the container manager installed on a device (e.g., Ecs container type) in order to obtain an application/service container on demand and to start using the startup parameters specified in the container descriptor.

6.6 Example Applications

For diving into the platform concepts and operations from a practical perspective, examples are desirable. One source for such examples are the configuration model regression tests, which contains two simple applications (SimpleMesh and SimpleMesh3) as well as a more complex example (RoutingTest). However, due to their very nature, these test cases are rather artificial and use a single implementation project for the realization of their services to ease the overall build process. While this is acceptable for regression tests, it may not be an adequate starting point to better understand the IIP-Ecosphere platform.

To make the platform accessible and understandable, we started to collect some examples that can be published without breaking IPR. On the one side, this set of examples is still small and in development (to cope with unavoidable platform changes). We aim at providing more examples and accompanying explanation (such as slides or videos) in the future. On the other side, the platform example set is one extension point, where you can easily hook into and contribute own examples (also discussing alternative approaches, e.g., to code organization).

It is important to mention that the examples are meant to be created/completed by the platform instantiation. In other words, when you import and open one of the examples the first time in an Integrated Development Environment or just by executing Maven, **you may run into compile errors or missing dependencies**. These are (usually) not bugs, rather than parts that were intentionally not committed and that must be created using the platform instantiation process, i.e., some of the command interaction points of the PlatformInstantiator.

We briefly summarize the individual examples below. The examples can be obtained from GitHub¹¹². Each example includes a README.md file with more detailed explanations. Typically, a Linux build script is included, which is used for regression testing of the platform. The example set consists of

- **examples.python:** This example focuses on the integration of an asynchronous service realized in Python (through the Python service environment) and demonstrates the integration with accompanying Java services (a simple data source as well as a simple asynchronous data sink). Please recall, the Python service implementations must be based on (generated) platform interfaces, located in a specific module (services) and packaged into an own Maven artifact (type “python”) to become available to the platform instantiation. Similarly, the Java services for sources and sinks must realize respective interfaces and be packaged. However, as in the VDW example, neither the actual configuration meta-model nor the required interfaces are available after download but must be obtained/created through instantiation steps. As the services are linked in a stream-based manner, also a broker must be available. For this purpose, using the command (`mvn -U install`) you obtain the configuration meta-model, instantiate the broker based on the example configuration, create the application interfaces and basic implementation, compile the example, and finally integrate the service code in the example with the app. The application can be executed through a running platform, but also standalone, which requires a specific setup that is included in the example.

¹¹² <https://github.com/iip-ecosphere/platform/tree/main/platform/examples>

- **examples.pythonSync:** This example follows the same setup as `examples.python`, but integrates the Python service code into a synchronous service passing its data on to a synchronous sink. The remaining properties of the example are the same.
- **examples.rtsa:** This example shows how to utilize the RapidMiner RTSA as service in an IIP-Ecosphere platform application. Regarding the setup, the example is rather similar to `examples.python`, except for the use of RTSA instead of a manually implemented Python service. One specific aspect is that RTSA is commercial software and not included in the example. As alternative, we use our simple FakeRTSA from the regression tests. As above, using the command (`mvn -U install`) you obtain the configuration meta-model, instantiate the broker based on the example configuration, create the application interfaces and basic implementation, compile the example, and finally integrate the service code in the example with the app. The application can be executed through a running platform, but also standalone, which requires a specific setup that is included in the example.
- **examples.KODEX:** This example illustrates how to use KIPROTECT Kodex as anonymization/pseudonymization service in an application. Again, the example is rather similar to the RTSA and the Python example. One specific aspect is that Kodex is a generic platform-provided service that is customized for the use in your application based on the application configuration, e.g., what data to anonymize. A further aspect is that this example integrates an in-memory database as data sink. The build steps are the same, i.e., using the command (`mvn -U install`). Also this application can be executed through a running platform, but also standalone, which requires a specific setup that is included in the example.
- **examples.template:** This example contains the configuration model and the service implementation used as blueprint for the service development workshop in November 2022 and for the service development tutorial videos⁶. The example structurally is similar to `examples.python` and can be built in the same manner (`mvn -U install`).
- **examples.hm22:** For the Hannover Messe 2022 (HM'22) and, in an improved version, for the Tage der Digitalen Technologie in Berlin (TddT'22) the IIP-Ecosphere project team developed a demonstrator application that involves IoT/Factory hardware, the IIP-Ecosphere platform as well as a controlling application. For short, the application demonstrates a visual AI-based quality inspection process for a configurable lot 1 production, where a robot with mounted camera takes pictures from three sides of an aluminum car model. A QR code on the indicates an (external) AAS describing the configuration of the car. The two other pictures are utilized by a Python-based AI to detect the number of windows, the color of the tires, the presence of an engraving as well as whether “scratches” were accidentally caused on the surface of the car. The services, in particular the AI, can be executed in different deployment settings, e.g., on a Phoenix Contact AXC 3152 PLC/edge device. For this purpose, the application consists of two generated connectors (OPC UA to obtain input from the PLC, AAS to request information on the car configuration), a camera source, the AI (as service family to switch between alternative AI implementations), the action decider (controlling the overall process) as well as a customized version of the generic, platform-provided TraceToAAS service. The inspection results are displayed on an Angular Web Application running on a tablet based on the information in the application AAS (via the TraceToAAS service). The application involves forward data flows as described above, but also backward data flows for controlling the process. Figure 69 illustrates the physical setup, more details on the application and the lessons learned while creating/integrating the application can be found in [11]. The build process is similar to the examples described above. The application can be executed standalone for testing as well as via the platform through deployment plans (examples included).

- examples.VDW:** The IIP-Ecosphere partner VDW is prominently involved in standardization activities for industrial production, in particular in OPC UA through the UMATI initiative. VDW offers a test server providing model instances according to OPC UA companion specs. As one of the connectors that is shipped with the IIP-Ecosphere platform is an OPC UA connector, an example based on that connector may be interesting to the reader. Connectors for the IIP-Ecosphere platform may be created by hand or, the preferred way, generated from a configuration model. This example illustrates both, a handcrafted connector as well as the integration of a generated connector into a demonstrating piece of code. It is important to mention that we do not read out the whole UMATI test server structure, rather than just a small piece for the OPC UA Woodworking Companion Spec¹¹³. The example represents the respective part of the Companion Spec in terms of IVML, imports the structure into the configuration model, adds some information about caching and server coordinates and creates a simple application mesh that just consists of a source (the connector), i.e., no further processing with the obtained information happens here (and some of the generated artifacts may be unusually empty). For executing this example, using the command (`mvn -U install`) you first execute build steps to obtain the actual configuration meta-model (intentionally not included), run the application generation and compile the example.
- examples.templates:** This example project displays the usage of `impl.model` to create an application with three services of which one is a python service utilizing AI. It consists of two parts, the `examples.templates.model` and the `examples.templates.impl`. The `model` part is based on the `impl.model` and contains the configuration `.ivml` files, with the use of `mvn generate-sources` we can generate templates fitting these configuration files. The `examples.templates.impl` is created by importing the resulting `.zip` file in the `gen` directory of the `examples.templates.model` project, it contains the implementation of the concrete service functionality. The services implemented in this project are made available to the `examples.templates.model` project by running `mvn install` in the `.impl` project. The application is completely build by then running `mvn install` in the `.model` part of the project. The result is an application that takes the data from the resource directory and evaluates it utilizing a python service with a pre-trained random forest classifier.



Figure 69: Using the IIP-Ecosphere platform to realize a robot-based visual quality inspection for Hannover Fair 2022.

¹¹³ Woodworking is not really related to the aims of IIP-Ecosphere, but it was the first one that we identified as potential candidate in the UMATI test server and that was reasonable large but also not too large to be turned into an IVML model (automated work in this direction is planned for the next release).

6.7 Creating an Application

Due to the configuration and instantiation process, the (manual) creation of an application for the platform is not just a matter of some programming tasks. In this section, we summarize and detail the steps, which in future versions shall be better supported by respective platform tooling. For this version, you shall be able to realize code in Java or Python (depending on the configuration that you will create) as well as get to know how to work with Maven. More details on project structures or default build sequences will be presented in the next sections.

Figure 70, summarizes all steps that are necessary to create an application consisting of Java services. We assume that you use `example.python` or `example.rtsa` as blueprint. While Figure 70 is meant to provide an overview, we briefly discuss now the individual steps (assuming that you have a most recent version of the platform sources from GitHub in your local workspace). For more details, please follow our tutorial videos on service implementation⁶.

- 1) Download the template project `impl.model` from GitHub¹¹⁴. `impl.model` contains a setup for the configuration and code generation. You may rename the projects as you desire. Akin, you may rename the maven artifact name in the respective `pom.xml` file. We will continue referring to the project as `impl.model`.
- 2) In `impl.model` you need to adjust the `.ivml` files in `/src/test/easy` to fit your needs i.e. adding or changing datatypes, for this you can refer to our example projects as a guidance i.e. `src/test/easy` in `examples.rtsa` or `examples.python`, the install package or `SimpleMesh` and `SimpleMesh3`. If you plan to have multiple examples/configurations on the same machine, please consider changing the Maven artifact name in the variable `sharedArtifact` in `TechnicalSetup.ivml` so that each project can be based on individual interface artifacts. Make sure your `pom.xml` does contain the properties `<iip.model>` with the value `PlatformConfiguration` as well as `<iip.resources>` with the value `resources`. Afterwards execute `mvn -U generate-sources` to obtain an actual copy of the IIP-Ecosphere configuration meta-model and create the `/gen/py/` directory containing `ApplicationInterfaces`, the application itself and a template directory with the `.zip` to implement the services. The `ApplicationInterfaces` are needed for the service creation and will be added to your maven repository. To utilize the templates, you we recommend importing them as a separate project right besides your `impl.model` project. The template project will contain template Java service sources, a matching Java jUnit test suite, template Python service sources, tests and a Python Eclipse editor setup file for PyDev, if configured, template mocking test input files, and an initial YAML-based identity store file.
- 3) Import the generated implementation template into your Eclipse (we refer to it as project `impl.impl`). Depending on the services stated in your model, now the services must be implemented in terms of Java or Python code, respectively. Please consider in particular architectural constraint C7 stating that generated code must not be altered. When the services are ready, compiled and tested, execute `mvn install` so that the artifacts of `impl.impl` are installed and become available. Also implement the unit tests for the services and validate your implementation. For later steps, adjust the mocking input JSON files in `src/test/resources`.
- 4) After successfully running `mvn install` in the template project, you can go back to your `impl.model` project and run `mvn install` there, make sure that the values for artifact in your `.ivml` configuration are correct and point to the location of the services from the template project. You can check this by looking at the `POM.xml` of the template model and reading the `groupId` as well as `artifactId`.

¹¹⁴ <https://github.com/iip-ecosphere/platform/tree/main/platform/tools>

- 5) Step 5 generates the full application artifact for Spring Cloud Stream. We recommend to execute generated Spring Cloud Stream test cases (documented in README.md of `impl.model`) for each individual service first. Then, you may execute the application on your local machine (as documented in README.md of `impl.model`) and finally on the platform (see also Sections 3.13 and 8.4).

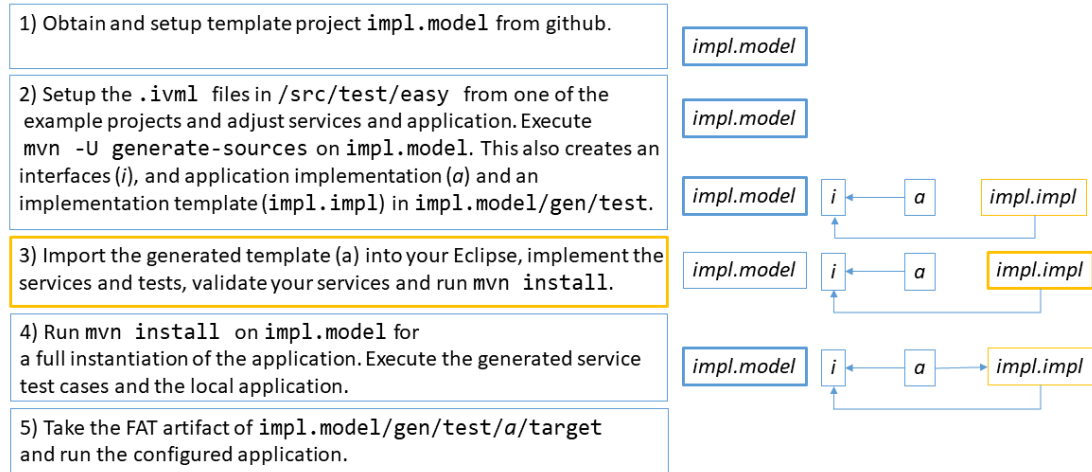


Figure 70: Steps to manually define services and create an application.

6.8 Project Structures

Due to the Maven build process and the generation of code during the instantiation, the service implementation creates/assumes a certain project structure that we will introduce in this section. As creating structures across multiple programming languages may be challenging, we also offer generated template projects for configured applications/services as we will discuss at the end of this section.

Figure 70 illustrates the overall structure of an all-in-one implementation project as we use it for most of the platform examples. On top-level, it consists of three folders, one folder for the generated code (`gen`), one folder for the own sources (`src`) relying on the generated code, the target folder containing the created/compiled binaries and the Maven build specification (`pom.xml`). There may be further files, e.g., for the integration into a CI environment.

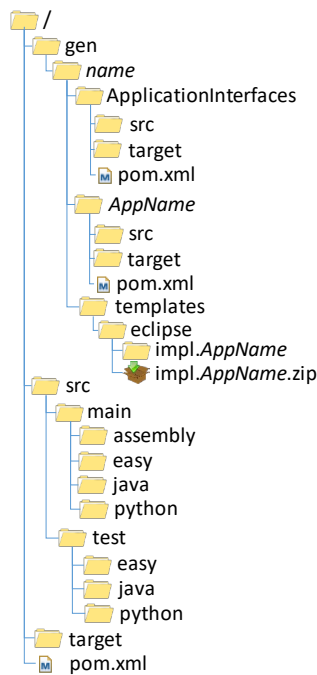


Figure 71: Overall structure of an implementation project.

Initially, the `gen` folder may not exist as it is created and filled during the instantiation of the project. This may lead to the effect that after an initial checkout, (local) dependencies are not in place, i.e., your IDE shows errors. After executing the instantiation/code generation and updating your implementation project, these errors usually disappear. The `src` folder contains your code based on the generated code. `src/main/assembly` contains project-specific Maven packaging scripts, e.g., for Python. These scripts must be called from the main `pom.xml` of the project. `src/main/easy` contains the platform configuration (meta-)model, which is downloaded as part of the build process. `src/main/java` contains production Java code, similarly `src/main/python` (may not exist if your services do not need python scripts). Similarly, the `src/test` tree contains testing code, e.g., for Java or Python. Moreover, as a “heritage” from the first examples, the `src/test/easy` folder currently contains the configuration model of your application (based on the model in `src/main/easy`).

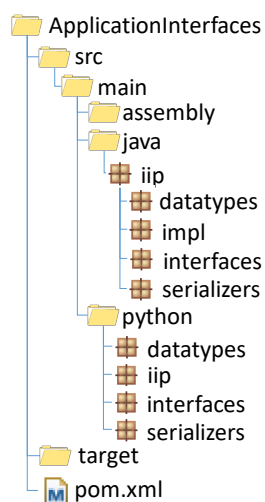


Figure 72: Detailed structure of the generated application interfaces.

After the first instantiation run, the `gen` folder consists of three sub-folders and a template POM file. The `ApplicationInterfaces` folder (if the name was not changed in the configuration model)

contains the data type interfaces, the data type serializers and the service interfaces for all applications of a platform installation. The second folder, which is named based on the respective application configuration in the model, contains the service integration into the service execution engine and service engine specific testing code, which relies on the `ApplicationInterfaces`, but also on your code. Thus, besides downloading the configuration (meta)-model, the build process consists of creating/deploying the sources in the `ApplicationInterfaces` folder, the compilation of your code based on `ApplicationInterfaces` and the generation of the second folder with the service integrations, which leads to the packaged application. The typical structure, which is rather similar to the overall project structure, is shown in Figure 72. The difference comes from the generated code and assembly descriptors and the assumptions about the underlying code structure. On the Java side, `datatypes` (representing the application data to be passed between connectors and services), the implementation (`impl`) of default services as basis for your service implementation, the interfaces for the services and the wire format serializers (based on the format selected in the configuration model) are generated. Similarly, on the Python side, `datatypes`, service interfaces and serializers are generated. The `iip` folder contains the Python service environment and is created/filled during the build process. If Python-based services are configured, respective assembly descriptors are generated and linked into the generated build process in `pom.xml`.

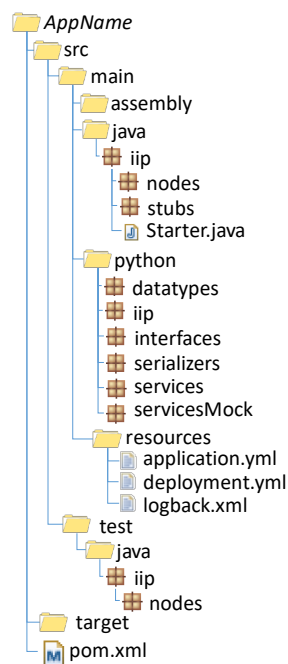


Figure 73: Detailed structure of the generated service integrations.

The name of the generated folder for the application depends on the configured application name. Also this folder represents a typical Maven project structure, here meant to complement the application interfaces and your implemented code in order to derive and package a full application. Similar to the application interfaces, assembly descriptors, Java code and, if configured, Python code is generated. The Java code consists of the nodes representing the integration of individual connectors or services into the configured service execution environment (by default Spring Cloud Stream). The stubs contain local parts of remote service implementations via AAS¹¹⁵. The `Starter` class is the main class of the application and responsible for serializer and service registration (as an extension of the respective class in the service environment). The `python` folder accumulates the generated code from the `ApplicationInterfaces` including the Python service environment (both parts

¹¹⁵ The stubs are currently not used and may be removed in a future release.

downloaded/updated during the build process). The `services` folder is the default place for Python service implementations and taken over from your implementation project (similarly the `servicesMock` folder, which contains Python service implementations for testing). The `resources` folder contains additional files that shall be made available on the classpath, in particular the Spring application configuration, the IIP-Ecosphere service deployment descriptor and the logging configuration. The `test` folder currently contains only a folder for Java sources, here the generated connector and service tests. The `target` folder contains the compiled binary and, in particular, the packaged applications, here in the formats discussed in Section 3.7.3.1.

The third directory contains the application/service implementation templates. This folder contains one folder per supported IDE (currently only Eclipse) and within that one folder for the template sources and one ZIP file containing the template sources to simplify the import into Eclipse. The template project contains a `README.md` detailing the generated parts and pieces, some top-level configuration files as well as a `src` folder for the production code (`main`) and the `tests`. Both folders may contain Java source code, Python source code and resources according to your configured services.

As mentioned above, most of the provided platform examples follow this structure, in particular to reduce the number of projects in an IDE. In addition, services may be implemented in two separate projects, one containing the model and one the service implementation. Templates for such a setup can be found in the GitHub repository of the platform.

For the future, we envision such a separation allowing the platform to take control over its configuration model. While it then shall still be possible to access the configuration model for download, it will not be in the control of the user/developer anymore. It is then a prerequisite that the platform offers means to modify the configuration model and to define new/modify existing applications. For this purpose, the configuration component shall mirror its configuration into the platform AAS and allow the management UI to access and modify the configuration. Ensuring the validity of the platform configuration, executing the instantiation process and providing access to implementation-level artifacts will then be in the responsibility of the platform. First steps in this direction are already taken in this version: an initial access to the configuration model through the AAS of the configuration component and the creation of template implementation projects (for Eclipse). A generated template project contains template files to be extended, completed by the developer, assuming for now that a respective `impl.model` project exists. A generated template project for Eclipse as illustrated in Figure 75 contains in particular a

- `README.md` explaining the build process and the test steps
- Eclipse specific files, i.e., `.project` and `.classpath`
- An initial `.gitignore` file
- A `pom.xml` Maven build process file containing the basic dependencies and linking the generated service/connector test cases.
- Java template source files according to the class names of Java services in the configuration file also considering the given packages of these class names.
- Python template service source files if Python services are configured. If Python services are configured, also the `python.xml` assembly descriptor will be generated and hooked into `pom.xml`.
- Test data specification in `src/test/resources` (to be filled and renamed)
- Initial identity store file in `src/main/resources`

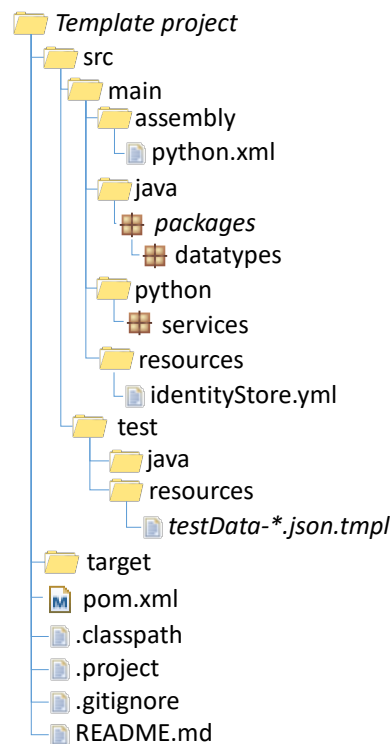


Figure 74: Detailed structure of a generated service implementation template.

6.9 Default Build Sequences

Section 6.7 provided an overview on how to implement applications with the platform, while Section 6.8 introduced usual implementation project structures. To ease the realization of different applications and their building, we suggest using a set of common build commands across the application implementation projects. These commands have been mentioned before and are summarized in this section.

The build process of the IIP-Ecosphere platform relies on Maven, i.e., is specified in Maven build specifications (`pom.xml`). The basis of a build specification based on the IIP-Ecosphere platform is the inclusion of the platform dependencies as parent project, which introduces common dependencies and their versions as well as common build steps and plugins (which can be refined in own build specifications if need be).

An application build specification consists of two types of build commands:

1. Build commands related to the configuration model and the (application/service) code generation provided by the platform instantiation process as discussed in Section 6.2. These commands involve obtaining the actual version of the configuration (meta-)model (`generate-sources`), instantiating the broker (`generateBroker`), generating interfaces (`generateAppsNoDeps`) and generating/integrating the full application (`generateApps`)
2. Build commands for the application/service code in the actual project, i.e., compiling the code against the generated application/service interfaces as well as tests of the (mocked) code in the actual project, usually the implemented services.

These two types involve different dependencies and define different build commands. In particular, it is important that some of the dependencies are not mixed among the two types as this usually leads to build errors. This difference is constituted by dependencies to artifacts created by the application/service instantiation, in particular `generateAppsNoDeps`. While the application/service code in the actual project requires the generated interfaces and the testing even the generated

services and service tests, the build commands for the configuration model do not require these dependencies as they are stated in the configuration model. Even worse, if these dependencies are declared for the build commands for the configuration model, they break the initial build as these dependencies are not present: These dependencies are created with the first run of `generateAppsNoDeps` and they are intentionally not deployed to the platform Maven repositories as they may contain IPR protected code, licensed components etc.

For this purpose, a typical build file defines two profiles, one profile per type of build commands. In addition, common dependencies may be defined in the (usual) main part of the Maven build specification, and, thus, are made available to both profiles. These profiles are:

1. **EasyGen:** This profile defined the build commands for configuration model and the code generation (the first type introduced above). As explained above, this profile must not have any dependencies on the code in the actual project and also not to Maven artifacts generated by the platform instantiation commands, in particular `generateAppsNoDeps`. By default, the commands declared in this profile call the `PlatformInstantiator` to create the `gen` folder and to instantiate the platform or application artifacts within there. For the IIP-Ecosphere platform examples, the commands are defined in a way that the application configuration model is in `src/test/easy` with an example-specific name and the platform configuration (meta-)model is stored in `src/main/easy`. The template projects include template build specification that can be adjusted to the naming/locations of your application/company.
2. **App:** This profile defines the specific dependencies of the application, in particular those that must not be declared in **EasyGen** as explained above (`ApplicationInterfaces` is usually a compile-scope dependency, the application specific artifact explained in Section 6.8 a test-scope dependency). Besides the build process for the project specific code, which typically is covered by the build steps defined in the platform dependencies, this profile may perform specific steps, e.g., packaging though artifact descriptors (mentioned in Section 6.8), but also executions of tests to ease the development (e.g., as in the template build project introduced in Section 6.8) or an execution of the full application for (mocked) dry runs. Usually, all these executions include the Maven test-scope and, thus, not only the `ApplicationInterfaces`, which are required for compilation, but also the application specific artifact containing the generated integration into the service execution environments as well as the generated connector/service tests (the latter also in test-scope).

In many cases it is sufficient to state the specifics of the build process in terms of static values, e.g., the folder where the application configuration model is located or the name of the IVML configuration (file). One particular exception is the resources folder. As discussed in Section 6.8, the resources folder contains files that shall be bundled with the application, including IPR-protected files. Currently, we follow the approach to define two folders, one with non-IPR content, e.g., a fake implementation of RapidMiner RTSA and a second folder with IPR content. The non-IPR content folder can be shared, e.g., in GitHub, while the IPR-content folder is not uploaded. For a correct instantiation, it is important to tell the instantiation process which folder to use and there are situations where you need instantiations for both folders. Thus, the recommended build sequence allows changing the resources folder from outside Maven, via `-Diip.resources="NewFolderName"`, where `NewFolderName` is the name of the actual resources folder.

6.10 Service Realization Rules and Considerations

Realizing an application within the conventions and assumptions of frameworks or an entire platform is not trivial, in particular if conventions and assumptions are not documented. While we detailed the architectural constraints for the platform in Section 4 and we will discuss the implementation in

Section 8 and frequently asked questions in Section 9, it is worth to think about service implementation considerations at this point, i.e., after the overview on creating an application in the previous section.

- As stated in Section 3.7.1, a **service is a long-running function** that may continuously be fed with data. A service starts when its status is set by the platform to **STARTING**. The responsibility of the service is to do then all preparations for starting, e.g., resource allocation, and to set itself into status **RUNNING** (or **FAILED** for good reasons). Akin to the startup, at any point in time, the platform may request a change to **STOPPING**, giving the service the opportunity to stop processing and to release resources gracefully. The responsibility of the service is to go to **STOPPED** when the cleanup is done.
- Both, startup and shutdown activities shall only **take time as long as absolutely necessary**. We are aware of the fact that some processes take a longer startup time, e.g., Python with TensorFlow or a complex Java service like RTSA.
- While a service is running, it is **kept in memory (to conceptually run forever)** by the platform in order to process data, i.e., it shall be ready to receive data and, depending on its function, produce output in synchronous or asynchronous manner. In other words, a service may keep its resources, ML models, etc. in memory. Of course, in particular on resource constrained devices, it makes sense to keep an eye on the overall memory consumption and to dynamically allocate/release resources that are not frequently used.
- A service **never returns/ingests** “the non-existing object”, e.g., `null` in Java or `None` in Python. Synchronous services always return a value upon their invocation, while asynchronous services may ingest any number of return values.
- **Take care for the executability**, i.e., do not kill your service/the executing JVM and catch/handle your exceptions properly. The reason for this is that the service execution immediately sets the **FAILED** state on a service that unexpectedly dies. This implies that you must also have an eye on the used libraries whether they cause such problems (JVM termination, throwing of runtime exceptions) and that you as service developer are responsible for handling such issues on behalf of your service.
- **Most state change activities are handled for you** in default implementations. In most cases, the existing/generated service frame allows you to hook into this process. In particular, platform supplied services do care for their own lifecycle state. Generated Python services are currently not informed about state changes. This will happen when extending the Python service environment. We expect no changes to existing service implementations as the Python services are based on classes with implementation, i.e., just some further methods will occur there.
- **Services are classes**. This is no big deal for Java as in Java every language unit is a class. However, it may be a surprise for Python, as there are also Python Scripts. The Python Service environment requires that a service is a class and that your operations are correctly declared with `self` parameters. The generated code tries to give you hints about the expected instances for input and output in terms of Python type annotations.
- **Services shall be re-entrant**, i.e., further input data may arrive while your service is processing. While synchronous services may be throttled by the service execution engine, for asynchronous services there is no external visible correlation between input and output data. In other words, for an asynchronous service, the service developer is responsible for handling parallel input data. It is also important to consider that the application design may prevent parallel inputs, but from the service side no assumptions on this can be made. Future versions of the platform may allow services to declare themselves as non-reentrant or a maximum parallelization degree.
- Please note further that **Java services may employ Python** functions, e.g., by executing a script. Here, the class requirement for Python does not apply, as the Java service calls Python

as a whole (in contrast to Python services that are hooked into the Python service environment). However, it is important that such Python functions are packaged correctly at compile/packaging time and unpackaged at runtime. For the compile time, a Maven assembly descriptor shall package all Python scripts and their required resources into a ZIP file and deploy it to the Maven repo (classifier `python`, type `zip`) along with accompanying Java code. The Java code shall use the `ProcessSupport` class of the Java service environment, which cares for correct unpacking and execution of Python scripts. The POM of the utilizing service must declare both, the Java and the Python artifact (classifier `python`, type `zip`) as dependencies.

- **Don't be blocking.** All service operations are expected to be executed as fast as possible, in particular service management operations. While there may be operations that take a certain amount of time, e.g., starting a JVM or a complex Python script within a service, longer running, resource consuming (blocking) operations are not permitted, e.g., reading a file and waiting for some reactions of the service. Such blocking operations may interfere with other services or the service control and leave the impression of timeouts, which may cause the respective service to enter the `FAILED` state.
- Do not make assumptions about **locations of file resources**, e.g., images or ML models. A service will be packaged (along with other services of the same application) into a service artifact. The layout of the artifact depends on the capabilities and conventions of the service execution engine. Currently, the platform only employs Spring Cloud Stream with two packaging strategies, see Section 3.7.3, and you do not know which artifact, i.e., which packaging layout, will ultimately be started by the user. Moreover, if further service engines occur, they may come with their own conventions. Thus, the following rules apply:
 - For **Java services**, please rely on the `ResourceLoader` (see Section 3.5.6), which is designed to cope with this insecurity in a systematic manner. If for some reason needed, you may introduce your own (local) resource resolver.
 - For **Python services**, please rely only on the existence of the files that you ship with your own service (see above), which is unpackaged into an own temporary directory at runtime. This directory is used as process working directory for the Python process, which then has access to all Python modules and resources, i.e., resources shall be loaded only through relative paths starting at the directory where the top-level Python modules of your service are located. Python Scripts used by Java services are treated similarly.

For any **other resource**, e.g., the temporary folder, do not make any assumptions about their location. While it is usually at `/tmp` under Linux, this may have changed during platform installation. For windows, the temporary folder is typically located somewhere in the user profile (partially also depending on Windows version and user interface language). Java, as well as Python offer programmatic access to the temporary folder, e.g., via `de.iip_ecosphere.platform.support.FileUtils` in Java or `tempfile.gettempdir()` in Python.

- Rely on **default service implementations** where feasible. For each service interface, we generate a service interface and a default implementation during the instantiation process (see also Section 6.2). For Java, interface and default implementation are separated, for Python the interface and implementation are formed by the same class. Currently, such a default implementation contains code for handling the service parameters as well as for the asynchronous data ingestors¹¹⁶. Due to type safety, both aspects, parameters and ingestors

¹¹⁶ An ingestor is a (lambda-)function that encapsulates the data ingestion so that the service developer does not have to care for stream names or routing aspects, which are typically specific for the selected service execution framework.

cannot be realized in a generic manner. Although it is not difficult to implement both aspects manually, it is also a tedious task. Moreover, it is a common programming error to miss adjusting the parameters or ingestors when your application model changes, e.g., when parameters are added or multiple ingestors are needed due to multiple output streams of a specific service. Inconsistent service implementations easily lead to long debugging tasks, while just a parameter declarator is missing, a parameter name is wrong or the service implementation expected to receive just a single ingestor. To ease the work of service developers and to keep up with changing models, we recommend to rely on generated default service implementations wherever feasible.

- Service output to the console is typically subject to **logging**. The logging target may depend on the active service execution engine, e.g., for Spring Cloud Stream, a temporary folder per service is created, which contains a folder with the deployment name and within that folder one file for the standard output or standard error stream. It is important to mention that depending on the use of Python for services, the logging target may differ. In the current version of the IIP-Ecosphere platform, Python services being executed through the Python service environment log into both streams, i.e., the log output appears separated. The process-based execution of Python functions, which may be used for service implementation, currently joins both streams for technical reasons and logs them to standard output.
- In addition to logging, further service output can be used for **debugging**, e.g., the generated (optional) logging or tracing of messages. Here, logging is more for local debugging while tracing (messages via the transport layer) can be helpful for remote debugging. To support remote debugging, the platform contains a transport message logger, which allows for receiving, emitting, storing and basic filtering of status, trace and monitoring messages.

7 Platform Security and Data Protection

In this section, we discuss means to ensure the security and the data protection in the IIP-Ecosphere platform. We start with (cross-cutting) internal security and security analysis in Section 7.1, the support/application of concepts of the IoT component profile in Section 7.2 and external security measures in Section 7.2.

7.1 Internal Security and Security/Privacy Analysis

One main step before managing security and offering security enhanced services is to review where in fact security is needed. Moreover, concerning the General Data Protection Regulation (GDPR), security and privacy aspects must be considered as early as possible in the design and development of a system (privacy and security by design principles). Architectural models, in fact, offer an excellent possibility to support the realization of privacy and security by design principles.

In Section 3.4.2 we introduced a UML profile called UMLsec. We further introduced two privacy checks secure links and secure dependency. Such checks provide a possibility to perform security and privacy checks on the design (architecture) of a platform provided using UML models.

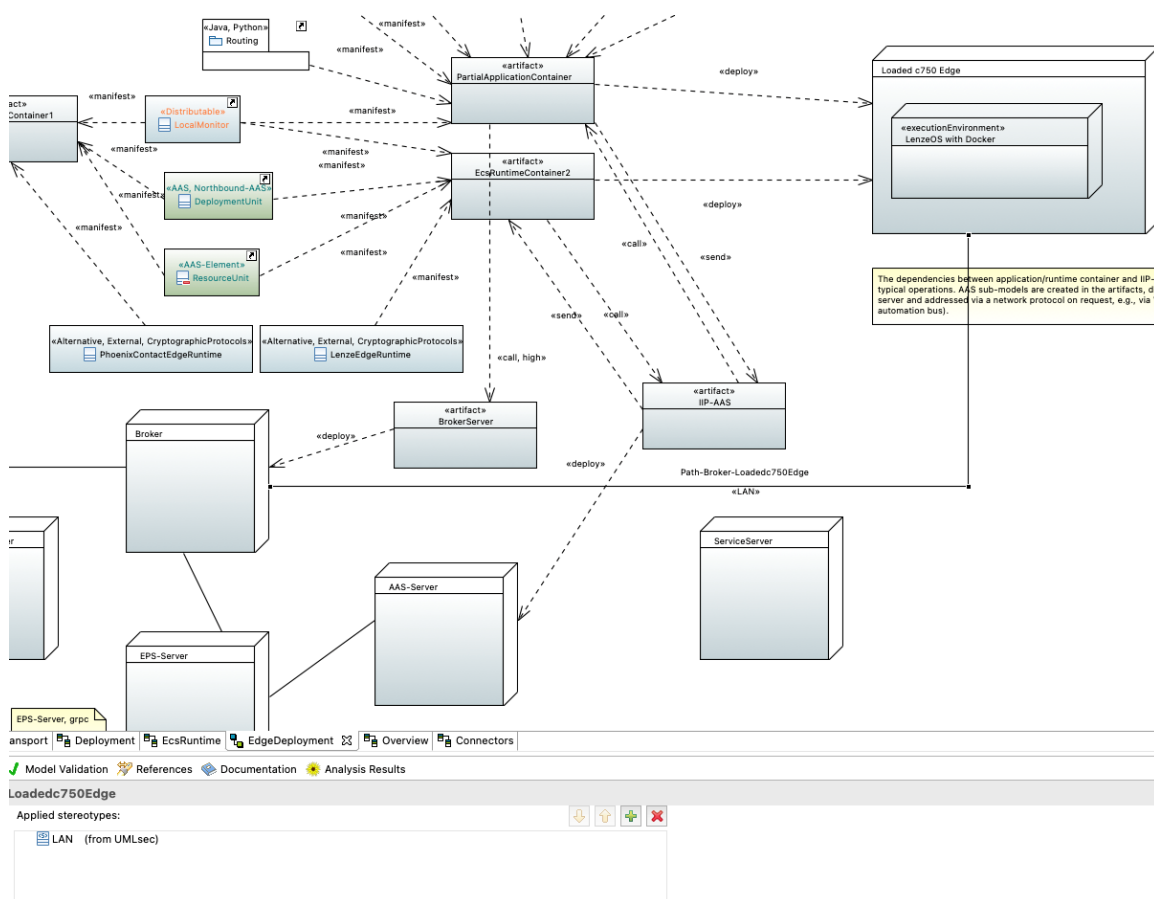


Figure 75: Architecture model for edge deployment annotated with secure links stereotypes (excerpt of Figure 42)

The process of checking and enhancing an architecture model is the topic of our ongoing research. In this section we describe how we can analyze the architecture of the IIP-Ecosphere data platform using CARISMA.

In Section 3.4.2.1, we saw that the *secure links* annotation enables one to ensure the security of communications in a physical layer. The following model is annotated with stereotypes relevant to secure links (Figure 75).

The link between the node Loaded c750 Edge and the node Broker in this figure is annotated with the Lan stereotype. The stereotypes as shown in the lower side of the figure can be set in properties view. In Figure 19 we saw that a default adversary cannot delete, read or insert on a link annotated with the Lan stereotype. Furthermore, the dependency between the two artifacts deployed on these two nodes namely, BrokerServer and PartialApplicationContainer is high (as indicated by the «high» annotation), requiring that the adversary cannot read, delete or insert on the link. Concerning the fact that the link is annotated with the Lan stereotype, after performing an analysis the check should not show any problems. This is in fact true, shown in Figure 76 which demonstrate the results of the analysis.

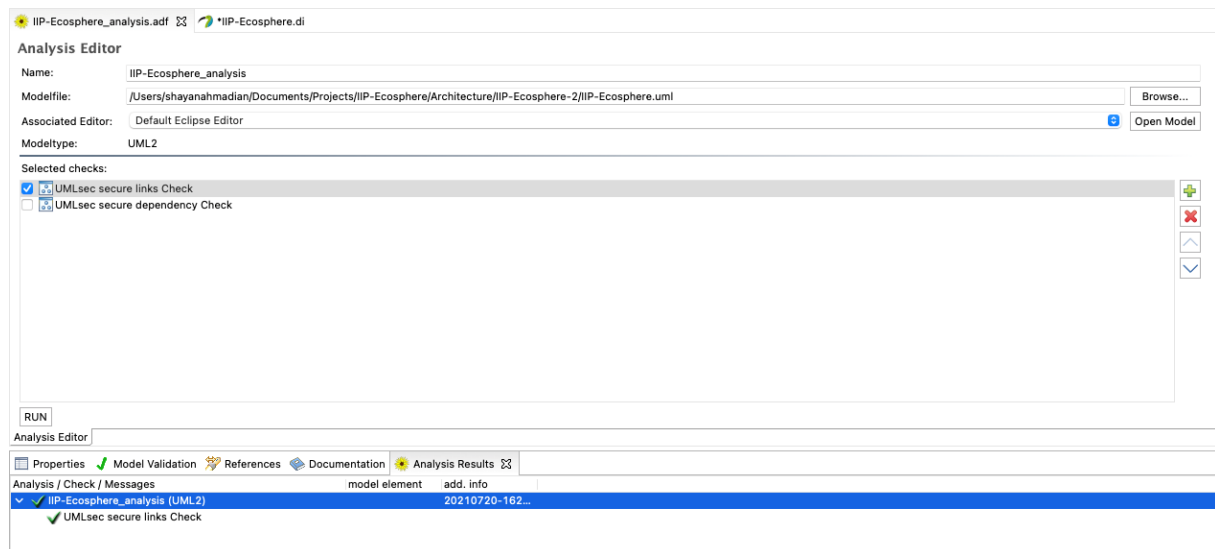


Figure 76: The CARiSMA analysis result.

Assuming that the link between the two nodes is annotated with the «internet» stereotype, and concerning the fact that the default adversary can read, insert and delete on the communication link, the corresponding error is shown in Figure 77.

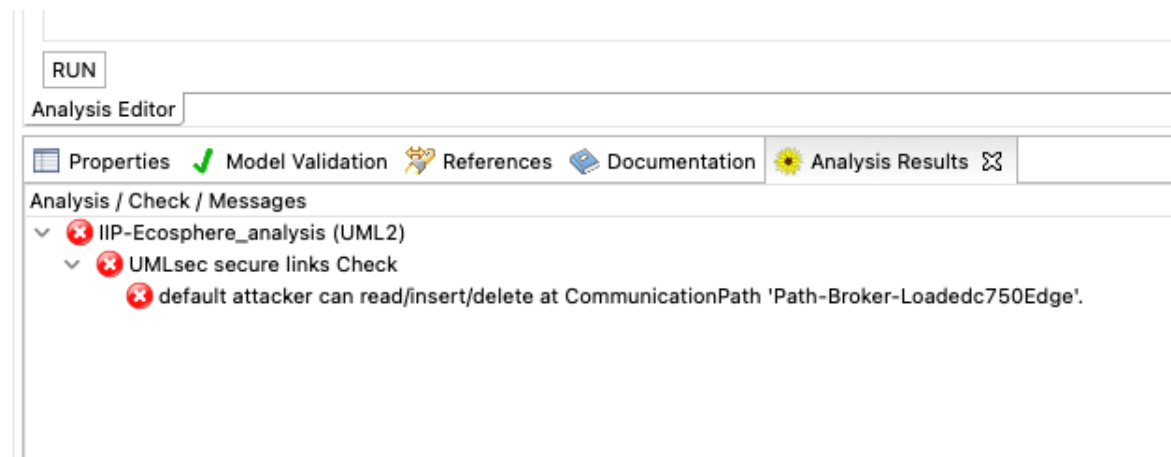


Figure 77: The result of the CARiSMA analysis.

If the communication path “Path-Broker-Loadedc750Edge” is annotated with the stereotype «Encrypted» and concerning the fact that a default adversary can only delete a message on an encrypted path, the corresponding error after performing a CARiSMA analysis is shown in Figure 78.

Such errors identified in the result of a CARiSMA analysis can inform a security expert or a system designer about potential threat and risks in a system in the very early phases of system design.

Therefore, such an analysis facilitates the process of enhancing an architectural model with appropriate security and privacy mechanisms.

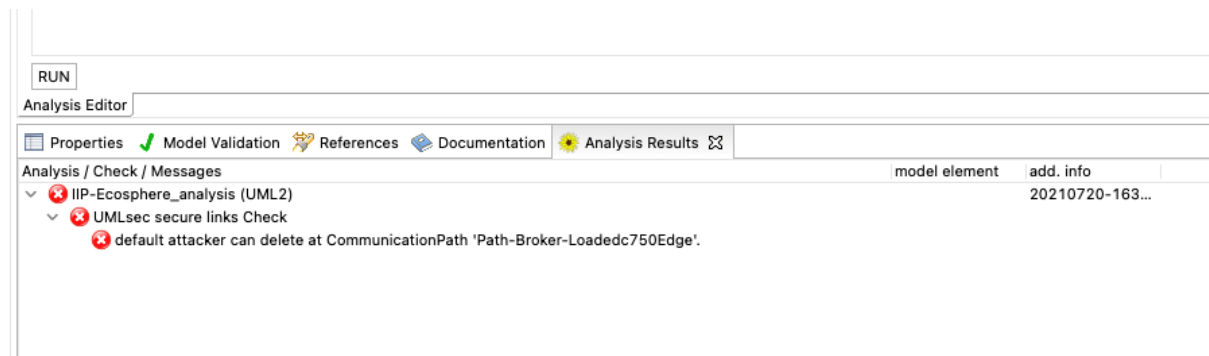


Figure 78: The result of the CARISMA analysis.

7.2 Support of the Concepts of the IoT Component Profile in actual Platforms

Technologies used to meet the protection goal of confidentiality are shown with the frequency with which the technologies are used in Figure 79. According to our studies the use of TLS to ensure confidentiality and integrity of data transmissions was common. Rarely was a VPN used to secure transmissions. Seldom was a way provided to encrypt stored data or to encrypt data by default. Responsibility over data confidentiality was often handed off to an external storage component. All of the platforms considered using Role Based Access Control (RBAC) as the authorization scheme for access control for users and components within the platform. A platform may also use Attribute Based Access Control (ABAC) in addition to RBAC.

Technologies used to meet the integrity protection goal are shown with the frequency with which the technologies are used in Figure 80. To ensure traceability, as part of the protection goals of integrity and transparency, some form of centralized logging was often used. Platforms also frequently have automatic patch management for devices at the edge. Even more common than automatic patch management for the edge is some form of fleet management. It is noticeable that all platforms that have automatic patch management for the edge also meet the requirements for fleet management. Within the platforms, two types of authentication were common. Devices at the edge mostly authenticate via certificates. In contrast, components within the platform itself often authenticate via a token mechanism. Less common are Http Basic Authentication, authentication via username/password, or to an identity provider.

Technologies that are used to fulfill the protection goal of availability are shown in Figure 81, along with the frequency with which the technologies are used. It was rarely mentioned explicitly that scaling of resources or the use of message buffers is used to maintain the availability of the platform. Only backups were found frequently enough to establish a trend, with six mentions. However, there is also the possibility here that responsibility for the availability of data is handed over to storage components outside the platform. However, this was not explicitly mentioned in any of the documentation.

Few, if any, technologies were identified to meet the protection goals of Unlinkability, Transparency, and Intervenability. Technologies that are used to fulfill the protection goal of Transparency are shown with the frequency with which the technologies are used in Figure 82. Rarely is data classified according to its protection need as it is collected. No technologies were identified to enforce Unlinkability that are not also used for Confidentiality. It is noticeable that the Intervenability objective has no associated technologies.

Some results from this study are provided in the following figures (Figure 79-Figure 83). They demonstrate the most used mechanisms to ensure the relating security or privacy goal.




Figure 79: Mechanisms used to ensure confidentiality.




Figure 80: Mechanisms used to ensure integrity.




Figure 81: Mechanisms used to ensure availability.




Figure 82: Mechanisms used to ensure transparency.

A special case was the way users authenticate within the platform. There was no technology that occurred frequently enough to identify a clear trend. However, all of the platforms studied have a defined way for users to authenticate. The different technologies are shown in Figure 83. It should be noted that OAuth 2.0 is the most commonly used technology in the twelve platforms, with only five mentions.

Platforms rarely defined a secure way to transport data out of the platform ecosystem. Based on these observations, it can be concluded that the focus in the development of the platforms studied was on the technologies classically assigned as part of IT security. In particular, these include the goals confidentiality, integrity and availability. Methods for enforcing data protection, on the other hand, were implemented rather rarely.




Figure 83: Mechanisms used to authenticate the users in IIoT ecosystems.

The feature model introduced in Figure 24 is in fact for two main purposes created. Firstly, it provides a structured means to organize the mechanisms and technologies that ensure security and privacy goals in IIoT platforms. Later this model can be iteratively expanded to include more mechanisms and technologies provided in a structured way. Secondly, it provides a basis to perform a model-based analysis to investigate if various components in a system model which represents an IIoT platform satisfy a required level of security or data protection.

7.2.1 Using the Profile to annotate a system model with proper mechanisms

In this section we describe the usage of the feature model to annotate UML models with privacy and security mechanisms. The feature model (see Figure 24) was therefore first transferred to a separate UML profile. To transfer the feature model into a UML profile, the navigation direction was inverted and replaced by an extension. Thus, a directed association from confidentiality to encryption became "Encryption extends Confidentiality". An example of how the translation between the feature model and the UML profile works is shown in Figure 84

The «Metaclass» class and «Metaclass» interface were specified as metaclasses for all targets in order to be able to attach all stereotypes to classes and interfaces. It is assumed that classes and interfaces can represent components of (I)IoT platforms. Since the feature model and thus also the profile lists a large number of technologies and implementations to protect the confidentiality and integrity of data during transport, the objectives confidentiality and integrity were also specified as the «Metaclass» association. It should also be possible to check the confidentiality and integrity of data in transit. The assignment of protection goals to metaclasses is shown in Figure 85.

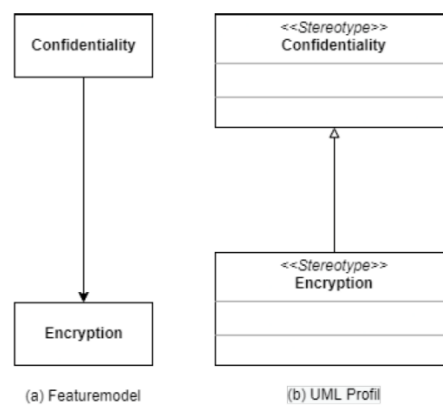


Figure 84: An example showing the realization of the feature model as a UML profile.

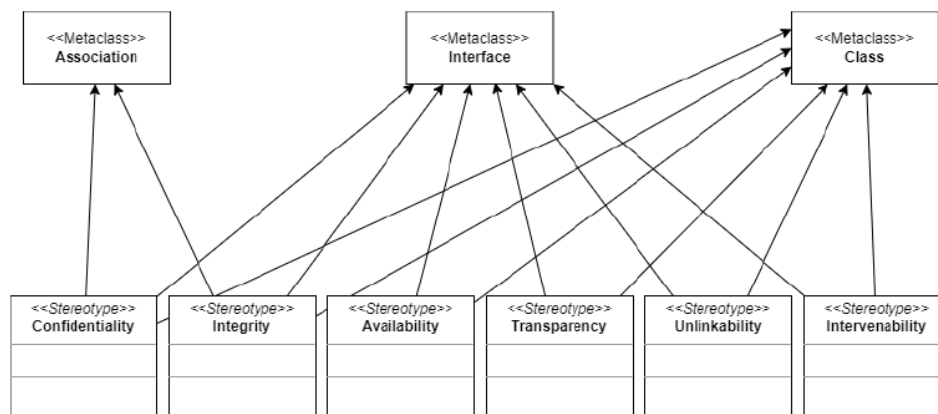


Figure 85: The assignment of the security/privacy goals to the meta classes in the UML profile.

7.2.2 Towards an automated analysis to verify required security levels

In this section we introduce rules that are formal enough to be implemented in automatic checks that can be used to verify the required security level of a system model.

Rule 1: A goal of a class is considered fulfilled if a class connected via an association has at least one stereotype that fulfills the goal. In this case, the second class is assumed to provide the first class with technology to fulfill the first class' goal. An example of this case is shown in Figure 86. See Figure 87 for an example that would not satisfy this rule. Only the stereotypes of classes that are directly linked to

the examined class via an association are considered. Classes that are connected to the examined class via more than one association and thus also via other classes are not considered, since these are not necessarily integrated into the functionality of the examined class.

Rule 2: A goal for a class is considered fulfilled if the class itself has at least one stereotype that fulfills the goal. In this case it is assumed that the class itself has technologies to fulfill the objective. A simple example is shown in Figure 88.

Rule 3: A goal of an association is considered fulfilled if both classes connected via the association have at least one stereotype that satisfies the goal. In this case, both classes are required to have the same stereotype, since the same technology must be used when both classes communicate via the annotated association. For example, in the case of communication, the association, between two components, the classes, it can be assumed that the communication is only successful if both components use the same protocol and the protocol represents a technology that meets the protection goal of the communication. An example of a UML model in which this rule is fulfilled is shown in Figure 89.

If all goals in the model are met, the analysis is successfully completed. The goals that have been met are listed and output. If not all goals are met, a search is made for solutions. As suggested solutions, for all elements that have unfulfilled protection goals, other elements in the model should be listed that have at least one stereotype with which the protection goal can be met and list which stereotype(s) fulfill the protection goal. After that, the analysis is terminated as unsuccessful. This process is shown in Figure 90.



Figure 86: An example for the rule 1.



Figure 87: An example where rule 1 is not satisfied.




Figure 88: An example where rule 2 is satisfied.



Figure 89: An example where rule 3 is satisfied.

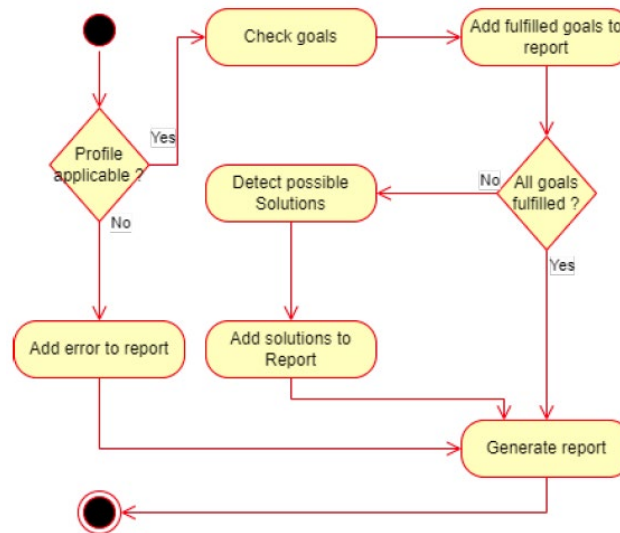


Figure 90: The workflow of the check.

In the following example we consider the Connectors View shown in Figure 91.

First, the protection goals are then added to the corresponding points in the model or view. For this demonstration the “Connectors” class is annotated with the <<Confidentiality>> and <<Integrity>> stereotypes.

This demonstration looks at the Eclipse Leshan and Eclipse Californium components. The excerpts of the templates is shown in Figure 93.

The result shows that not all protection goals in the model are met. The result is shown in Figure 93. The detailed report shows that neither confidentiality nor integrity on connectors is fulfilled. As a suggestion to meet the Integrity on <<HMAC>> from the Eclipse Californium and Eclipse Leshan classes. The <<encryption>> stereotype from the Eclipse Californium and Eclipse Leshan classes is suggested as a suggestion for fulfilling confidentiality. In order to fix the detected errors after the failed analysis, the classes Lwm2mConnector and CoapConnector are annotated with the stereotypes <<encryption, HMAC>>.

We can perform this annotation of stereotypes here because the function of the Lwm2mConnector and CoapConnector classes is implemented at this point in the model by the Eclipse Californium and Eclipse Leshan classes. Therefore, at this point, the stereotypes of the Eclipse Californium and Eclipse Leshan classes are inherited by the Lwm2mConnector and CoapConnector classes. If after the classes Lwm2mConnector and CoapConnector were annotated with <<encryption, HMAC>> the analysis is executed again. The result of the second analysis is a successful check and is shown in Figure 94. The stereotypes <<encryption, HMAC>> for the two classes Eclipse Californium and Eclipse Leshan are derived from the documentation of the external components.

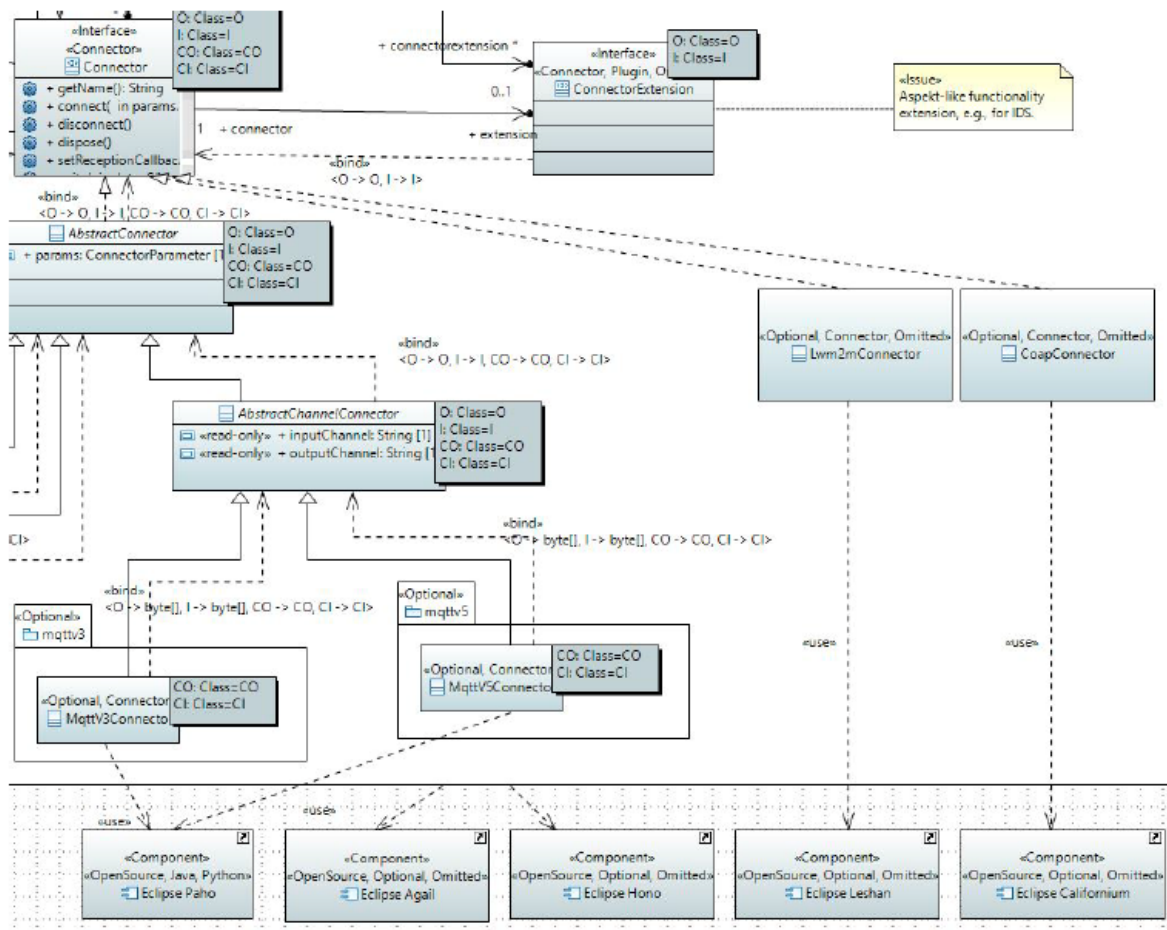


Figure 91: An excerpt of IIP-Ecosphere system model demonstrating the Connectors view.

Name	Eclipse Leshan
Documentation	https://github.com/eclipse/leshan/wiki
Wie wird Vertraulichkeit der Übertragung sichergestellt Schlagworte : TLS, AMQP, Modbus, MQTT, HTTPS, IPSec, VPN	Daten werden durch DTLS verschlüsselt
Wie wird Authentizität der Übertragung sichergestellt ? Schlagworte : TLS, Signatur, Zertifikat, X.509	Unterstützung für X.509 Zertifikate
Wie wird Integrität der Übertragung sichergestellt Schlagworte : TLS, HMAC, MAC	DTLS stellt Integrität durch HMAC sicher
Wie wird Vertraulichkeit von gespeicherten Daten sichergestellt ? Schlagworte : Verschlüsselung, AES	N/A

Figure 92: An excerpt of a table showing the results of investigating the security/privacy mechanisms and technologies provided by Eclipse Leshan.

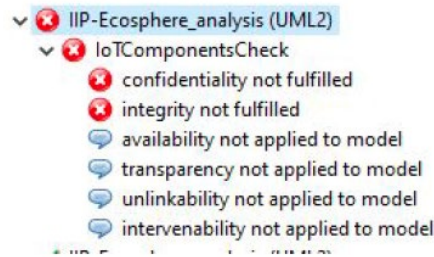


Figure 93: The results of analyzing the IIP-Ecosphere model. The confidentiality and integrity are not satisfied.




Figure 94: The results of analyzing the IIP-Ecosphere model. The confidentiality and integrity are satisfied.

7.3 External Security

For communicating with other platforms or other instances of the IIP-Ecosphere platform, external communication is required. In particular, external communication requires a certain level of security. Currently, two approaches are intended to support such external communication:

- Concepts and components of the International Data Spaces (IDS) that will partly be integrated into the GAIA-X initiative. We will discuss the plans for IDS support in one of the next versions of this document.
- An alternative, more lightweight approach is to transparently encrypt all communication between two parties linked via the internet. KIPROTECT has demonstrated such an approach in terms of the EPS-System (End Point Server, Figure 95 and Figure 96), which is already being used to secure data transfers between German health departments and contact tracing providers. In particular, it is easy to install, offers end to end encryption and mutual authentication via mTLS and supports encapsulation of arbitrary application-layer protocols like MQTT or REST interfaces. This may allow for a secure, transparent communication between two parties, e.g., two installations of the IIP-Ecosphere platform. Moreover, the EPS-System supports role-based access management that can restrict access to specific services and methods based on group memberships as well as additional criteria, so that the platform communication can be provided in a selective manner to actors in the ecosystem. Providing an additional level of authentication and access control via the EPS system on top of the existing ones of the platform can be part of a good “defense in depth” strategy.

Since version 1.3.0, BaSyx provides RBAC as one mechanism to protect the managed AAS, in particular their REST API. So far, the IIP-Ecosphere platform and its management UI only rely on TLS rather than authentication and RBAC, but we consider an upgrade of the security strategy for one of the next releases of the IIP-Ecosphere platform, which then requires exchange of access tokens (and potentially certificates) between the platform services, the device services and the management UI.

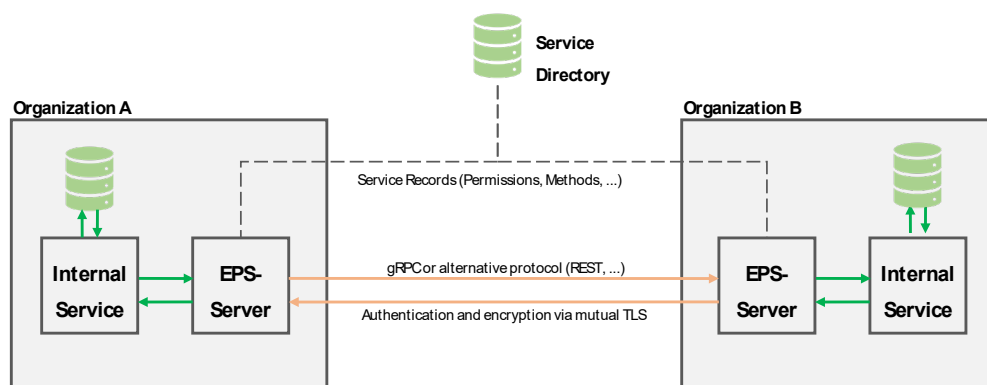


Figure 95: Direct communication between organizations through the EPS system




Figure 96: Indirect, proxy-based (end-to-end encrypted) communication between organizations through the EPS system.

8 Implementation

In this section, we briefly discuss aspects of the implementation of the IIP-Ecosphere platform, i.e., decisions we made during the implementation (Section 8.1), how to obtain a binary version (Section 8.2), the dependencies and how to compile the sources (Section 0), and how to install and to use the platform (Section 8.4). Section 9 on how-to's will take up some of the topics, but more from the perspective on how to extend or interact more deeply with the platform (code). Intentionally, we do not discuss code here. For this purpose, we refer the reader to the IIP-Ecosphere GitHub repository²¹ and in particular the Markdown¹¹⁷ readme files that are provided for the platform and for individual components.

8.1 Implementation decisions

We briefly discuss now technical decisions or issues that occurred during the development of the IIP-Ecosphere platform. This list may not be complete¹¹⁸ and is subject to incremental extension:

- As more parts and pieces show up, e.g., AAS sub-models, the more decisions on the **startup process** of the platform have to be made. However, some of these decisions impact testing, as a full startup including AAS sub-models is not always desired or may even break tests. In these cases, it is possible to mock out the AasFactory or to create missing server instances for the platform AAS via the AasPartRegistry, both located in the Support Layer (Section 3.5).
- Akin, many components make assumptions on **default instances for alternative components in testing**. Typically, we use AMQP as testing protocol (the server is rather easy to use and the implementation is stable) as well as BaSyx as AAS implementation. See architecture rule C5 and C6 in this respect.
- BaSyx and Spring use different versions of the **expression language** javax.el.el-api, which, when utilized together on the same classpath, prevent Spring Cloud Stream from starting. Wherever possible in installation packages, we try to separate AAS and stream processing, i.e., stream processing components shall run in their own JVMs controlled by a supervisor JVM containing the ECS runtime, which also maintains the representing AAS of the installation part. For uniform technical configuration, it is desirable that the ECS runtime is also started as a Spring application, while use of Spring Cloud Stream shall be prevented in there.
- Different external components depend on **Google Guava** in several versions. As a Guava version below 22 prevents some protocol test cases to be executed, we decided to fix Google Guava to version 22 in the platform dependency management. Similarly, further components may be fixed to rather narrow version ranges in the managed platform dependencies (and transitively in components such as EASy-Producer).
- So far, we use org.slf4j for **logging**, as this library is also used by BaSyx and Spring (although in different versions). Logging setup (also called configuration) is typically added during platform instantiation or for testing, also to avoid conflicting setups. In the logging setup as well as in the generated command line scripts we take care of actual security vulnerabilities and disable affected functionality. However, it is not possible to replace all potentially vulnerable versions, e.g., of logging frameworks, as partially the versions are determined by dependencies that we rely on and at the point in time of writing upgrading these components to more recent versions was considered as a project risk.
- We added a simple resilience mechanism for **failing connections** to AAS implementation servers. In the version of BaSyx that we are using, implementations of operations, property

¹¹⁷ <https://de.wikipedia.org/wiki/Markdown>

¹¹⁸ We do not intend to repeat all coding conventions for the platform in this document. We just listed here the most important ones with their rationales as overview. For details, please refer to <https://github.com/iip-ecosphere/platform/blob/main/platform/documentation/README.md>

getters or setters are attached through functors (usually lambda functions) to the AAS. In such a functor, currently the preferred style seems to be to create one connector instance per operation or property call, which builds up a network connection to an AAS implementation server. If the connection fails, e.g., because the AAS implementation server was intentionally shut down, let's say when stopping a service through the service manager, the AAS will continue connecting unsuccessfully to that AAS implementation server. At a glance, this only is an issue if the operation or property is addressed. However, in the BaSyx version used in this release, each access to a remotely deployed AAS causes an execution of all these functors (probably to serialize and transport the respective "value"), leading in some cases to (seemingly) endless trials to connect to the intentionally closed server. Although we delete the respective operation/property from the AAS before shutting down the service or the respective AAS implementation server, respectively, the described behavior occurs. As a mitigation and a first step towards **connection resilience**, the functors attached by the AAS abstraction for BaSyx track erroneous connections for all connector instances and return a constant value on failures. As this decision is intentionally global for all connector instances, we also have to revert the decision if the server becomes available again or the same address is used in another context. Currently, erroneous connections are disabled by default for a time period of one minute. Later versions of the platform shall integrate this behavior with the port release of the network manager or with a connection trial after a given timeout.

- As an operationalization of architectural constraint C9, we limit the type of the **exceptions** to be used in Java code. While in Java there is a style of throwing exceptions of different types for different purposes, we decided to use only a few types for a more uniform exception handling, namely, as checked exceptions `IOException` if any form of input/output may fail, `ExecutionException` if processing may fail (in particular indicating failing AAS operations) and `IllegalArgumentException` as the only unchecked exception if parameters are semantically wrong (e.g., for object construction)¹¹⁹. However, unchecked exceptions may be false friends and shall be used carefully. As discussed for C9, emitting exception information like stack traces to the console is generally forbidden and logging shall not be considered as the best or only option but rather as the ultimate option if no other exception handling like alternative/default processing or recovery is possible. Logging of exception traces shall only be applied if one cannot clearly distinguish between an error or a warning situation.
- Test artifacts shall be strictly separated from production artifacts as already mentioned in Section 8.1.

For easing the implementation of new components or examples, we provide several templates that already have built in conventions and development setup. Figure 97 illustrates the template for a Maven-based Java component¹²⁰.

- The name of the project indicates the component name that shall be reflected in the java packages for production and testing code as well as in the Maven POM file. The name of the test package shall start with "test", but if too many methods have to be declared public to be accessible to tests, the package may directly start with "de" to allow for package access. Packages shall be documented (package-info.java), the main test suite is `AllTests.java` in the test package (only tests declared there will be executed).
- The component is by default developed for JDK 11. If a more modern JDK for central platform components is needed, change the system library and the compiler settings in the Maven POM. Please take the smallest possible JDK version for compatibility, e.g., JDK 13.

¹¹⁹ Generally, please avoid such Ninja exceptions or at least document them where absolutely required in terms of (unneeded) throws clauses as well as in the related Javadoc comment.

¹²⁰ Located in <https://github.com/iip-ecosphere/platform/tree/main/platform/tools>

- The component template ships with a checkstyle setup taking the style information from the platformDependencies project (which must be obtained from GitHub and installed by Maven before) use.
- There is a default .gitignore file that excludes the target folder. Please ensure that only needed files that are not generated/obtained during the build process are committed, e.g., the configuration model (usually in src/main/easy) shall not be committed rather than obtained from its Maven deployment (usually phase generate-sources) so that model updates can easily be followed by issuing that command, also during automated builds.
- The build-jk.xml file is needed as starting point for the continuous integration (Jenkins, therefore “-jk”). This file refers to further ANT imports containing the settings for Jenkins as well as macros for Maven execution and deployment. As these files are in other repositories, it may be that your IDE issues errors about missing files. You can ignore these errors. Upon first use of the template, please change the name of the project as well as the pattern for the files to be taken from the Maven target folder for Maven deployment. In specific cases, further adjustments are needed here. Some example projects do not have this file rather than a build.sh Linux script, which in these cases is (so far) more convenient than an ANT file.
- pom.xml is the Maven build specification. Usually, it declares the IIP-Ecosphere platform dependencies as parent, defines only its artifactId (taking over the IIP-Ecosphere group), its deployment form, name, description and dependencies. The build plugins are inherited from the parent. In some cases, e.g., for obtaining/unpacking specific artifacts like the configuration model further steps can be added, which usually extend the existing build setup.
- Ultimately, README.md is the readable documentation of this component. It shall briefly explain the aim of the component, its setup (Yaml-Structure), its specific requirements/limitations but also actual issues and problems. Please keep this file up to date. Upon first commit, this file shall be linked into the parent GitHub folder README.md (see Figure 2 and the platform layering for selecting a proper folder) as well as in the top-level README.md file of the platform by HTML links.

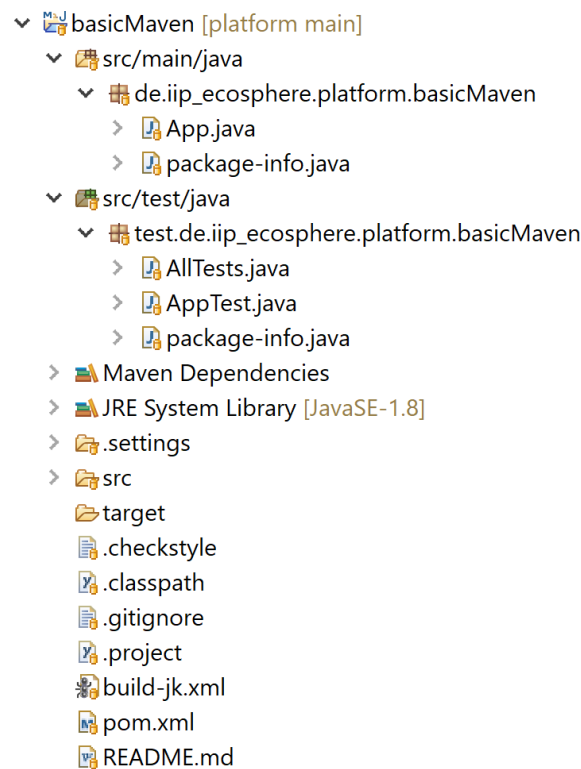


Figure 97: Structure of the component template “basicMaven” in the GitHub repository.

Along with this project, the `tools` folder of GitHub contains also the template project `impl.model` for the setup of applications. In contrast to the example structures, which unify both sides currently through two separate Maven files, these projects shall help setting up applications in separate folders. Moreover, the `tools` folder also contains the project `MavenCentral`, which contains tools for creating a release, like changing Maven version numbers, but also a tool to clean up a local Maven repository from superfluous SNAPSHOT versions.

8.2 Obtaining the IIP-Ecosphere platform

The sources of the IIP Ecosphere platform are available on GitHub¹²¹. Released binaries of the IIP-Ecosphere platform can be obtained from Maven Central¹²². Snapshots from the continuous integration can be obtained from the SSE Maven repository¹²³.

However, it is important to keep in mind that the IIP-Ecosphere platform consists of several alternative or optional components that must be consistently configured to obtain a valid installation for a certain setting. We will discuss in Section 8.4 how to utilize the configuration approach to obtain the binaries. Below, we summarize the (optional/alternative) components, the respective location of the configuration settings and the JSL descriptors that can be used to provide extensions. Table 24 summarizes the settings and the provided descriptors. Most descriptors define exactly one instance to be created, while descriptors indicated with * in Table 24 may specify multiple instances.

Table 24: Configuration and extension mechanisms used in the IIP-Ecosphere platform components (for descriptors, we abbreviate “`de.iip_ecosphere.platform`” by “`d.i.p`” for formatting reasons, * indicates the multiplicity)

Layer/ Component	Settings	Supported JSL descriptors
Support	-	<code>d.i.p.support.LifecycleDescriptor*</code> <code>d.i.p.support.LifecycleProfile*</code> <code>d.i.p.support.aas.AasFactoryDescriptor</code> <code>d.i.p.support.aas.ProtocolDescriptor</code> <code>d.i.p.support.aas.AasServerRecipeDescriptor</code> <code>d.i.p.support.net.NetworkManagerDescriptor</code> <code>d.i.p.support.iip_aas.AasContributor*</code> <code>d.i.p.support.iip_aas.IdProviderDescriptor</code> <code>d.i.p.support.metrics.SystemMetricsDescriptor*</code> <code>d.i.p.support.identities.IdentityStoreDescriptor</code> <code>d.i.p.support.resources.ResourceResolver*</code> <code>d.i.p.support.semanticId.SemanticIdResolverDescriptor*</code>
Transport	-	<code>d.i.p.transport.TransportFactoryDescriptor</code>
Connectors	-	<code>d.i.p.connectors.ConnectorDescriptor</code>
Services	<code>iipecosphere.yml</code>	<code>d.i.p.services.ServiceFactoryDescriptor</code>
Resources / Monitoring	<code>iipecosphere.yml</code>	<code>d.i.p.ecsRuntime.EcsFactoryDescriptor</code> <code>d.i.p.ecsRuntime.deviceAas.DeviceAasProviderDescriptor*</code>
Device management	<code>iipecosphere.yml</code>	<code>d.i.p.deviceMgt.storage.StorageFactoryDescriptor</code> <code>d.i.p.deviceMgt.registry.DeviceRegistryFactoryDescriptor</code>
Configuration	<code>iipecosphere.yml</code>	-
Platform	<code>iipecosphere.yml</code>	-

¹²¹ <https://github.com/iip-ecosphere/platform/>

¹²² <https://repo1.maven.org/maven2/de/iip-ecosphere/platform/>

¹²³ <https://projects.sse.uni-hildesheim.de/qm/maven/de/iip-ecosphere/platform/>

The *Support Component* does not take specific settings into account rather being set up through upper platform layers/components. In contrast, the Support component defines several fundamental JSL descriptors to allow the upstream platform components to hook into at defined points or to allow for external extensions. We summarize the descriptors now and link them to the variability provided by the platform and the platform configuration approach. The descriptors are:

- **LifecycleDescriptor** with allows adding components to the startup/shutdown process of a platform component. These descriptors can indicate a certain startup level and they even can cause a shutdown of a platform component. Adding certain descriptors to a platform binary causes the respective components to be started. Upper platform components ship with their descriptor file (in META-INF/services) so that either we add a certain component or a descriptor to the instantiated platform binaries (positive variability) becomes then active. In seldom cases, we may add the component and remove the descriptor to disable the respective registration (negative variability).
- As discussed in Section 3.5.3, lifecycle profiles may be a resort when dependencies run into conflict or individual parts of larger lifecycles shall be executed separately, e.g., to virtualize these parts. A **LifecycleProfile** allows for defining the lifecycle descriptor to be executed when the profile becomes active. A component may define multiple (alternative) profiles.
- **AasFactoryDescriptor** indicating the AAS factory to be used. A specific descriptor is shipped with the AAS (abstraction) implementation. The default implementation is `support.aas.basyx`. The platform just takes the first available descriptor (excluding potential descriptors used in testing), allowing here only for a single choice variability. By including a certain AAS implementation component, i.e., adding it to the platform classpath, the descriptor is made available and the respective factory becomes active (positive variability) as done during platform instantiation.
- **ProtocolDescriptor** is an optional extension descriptor indicating AAS implementation protocols that are not shipped with the platform. By default, TCP and HTTP/REST protocols for the BaSyx Virtual Automation Bus are provided, but other protocols may be desired in a certain installation. Here, additional external components can add arbitrary protocols (positive, unlimited variability) as long as the protocol names are unique. New protocols must be added to the configuration model as potential alternative so that the selected/desired protocol can be specified while instantiating the settings of the upper components.
- **AasServerRecipeDescriptor** defines the specific recipe to be used when creating an AAS server. The AAS abstraction defines a local server recipe for in-memory storage. However, on a server sided installation, also a persistent storage of the AAS may be required, which can lead to a large set of dependencies and unnecessary allocation of resources on edge devices. The required behavior, storage options but also dependencies can be defined by a specific AAS server component.
- **NetworkManagerDescriptor** is an optional descriptor that indicates which network manager shall be used by a component. The Support component does not ship with any descriptor information so that the platform instantiation must provide respective files (in META-INF/services). One alternative to the local manager is a global AAS-based network manager for some ports backed by a local network manager.
- **AasContributor** is an extension descriptor for higher level platform components to conveniently build a common AAS for the platform. AAS contributors define specific sub-models and announce their presence through the **AasContributor** descriptor, i.e., any descriptor found will be used to set up the common AAS for the platform. Examples in the Support Component are the platform “nameplate” sub-model and the network manager AAS (providing access to the network manager selected by the **NetworkManagerDescriptor**).

Although the descriptors shipped with the platform are intended to be executed, the platform instantiation may add or remove specific contributors to customize the AAS of a certain component. Moreover, the descriptors can declare themselves as invalid as, e.g., instances required to implement the AAS are not present.

- The `IdProviderDescriptor` is an optional descriptor indicating which strategy shall be used to determine the identification of a device. By default, the IIP-Ecosphere platform uses the MAC address of the device that is enumerated first by the system as device id. This default provider also allows for overriding the device id from via a command line parameter (cf. Table 28), in particular for testing. When used in an AAS, the provided device id may be modified so that it complies with the rules of an AAS short identifier, e.g., additional characters may be prefixed or whitespaces may be removed.
- The `SystemMetricsDescriptor` is an optional descriptor allowing to define a target-system specific implementation of the system metrics (cf. Section 3.5.4). Multiple system metrics descriptors can be specified, which may include system specific metrics descriptors. Of course, only one descriptor shall be active. Thus, with higher precedence, the first descriptor in the loading sequence is selected that declares itself as enabled, e.g., based on certain vendor system properties that uniquely identify a system installation. If none is enabled, a fallback descriptor, such as the default metrics via `jSensors` is activated. If no such fallback descriptor is available, the default metrics plugin is activated, which just returns constant values.
- The `IdentityStoreDescriptor` is an optional descriptor to define the actual implementation of the identity store. The default is the `YamlIdentityStore` reading `identityStore.yml`.
- The `ResourceResolver` realizes a strategy to load file-based program resources (cf. Section 3.5.6). The basic strategy loads from the root folder of a Jave archive. Additional resolvers may, e.g., treat sub-folders of such an archive as root folder for the resolution, e.g., as the resolver for the Spring FAT Jar format. As the `ResourceResolver` is descriptor and resolution strategy, we decided to drop the usual “Descriptor” name suffix here.
- The `SemanticIdResolverDescriptor` allows hooking one or multiple semantic id resolvers (see Section 3.5.7) into the platform. Two pre-defined descriptors are based on local excerpts of ECLASS and the `admin-shell.io` IRI namespace, a further one is the REST-based semantic id resolver utilizing the ECLASS web service.

The *Transport Component* in the Transport and Connectors layer does not define own settings. This is done by the `TransportFactoryDescriptor` to allow concrete transport protocol implementations to hook themselves into the `TransportFactory`. Moreover, it offers adding serializer implementations to the `SerializerRegistry`. For a concrete application, the respective serializers are created during platform instantiation and registered through generated code in the `SerializerRegistry`.

The *Connectors Component* in the Transport and Connectors layer defines the `ConnectorDescriptor` for announcing available descriptors that can be used / shown up in the AAS sub-model of the Connectors Component.

The *Services Component* takes setup information from a unified YAML file called `iiecosphere.yml`, which must be present on the classpath of the component¹²⁴. This file is instantiated through the configuration model during platform instantiation and added to the respective instantiated components. Depending on the service manager to be used, specific setup information may be required, e.g., for Spring Cloud Streams the full breath of the used Spring Components can be

¹²⁴ The detailed settings are documented in the `README.MD` file of the respective components.

configured in this file¹²⁵). Moreover, the Services component defines the `ServiceFactoryDescriptor` that announces the actual `ServiceManager` to be used.

Similarly, the *ECS Runtime Component* in the Resources and Monitoring Layer utilizes own entries in its `iiecosphere.yml` file and provides an own descriptor (`EcsFactoryDescriptor`) to announce the configured container manager. The AAS structure of the platform relies on devices with own AAS. However, it is unclear where the AAS for a certain device does come from, in particular if the device is already older or the vendor does not provide an AAS. To handle these cases, the ECS runtime allows to customize the AAS origin via the `DeviceAasProviderDescriptor`, which determines the component that returns the address of the respective device AAS (the component may also create the AAS if needed). By default, the platform reads/constructs an AAS from a simple Yaml file (with associated images) or an AASX file resolved from the classpath (see also Section 3.8.1).

In the Configuration Layer, the *Configuration Component* considers specific settings in its `iiecosphere.yml` file, e.g., where to find the configuration meta-model, the platform configuration, where to write instantiated components to etc. As the configuration component will offer own operations to modify the configuration, it also utilizes the descriptors defined by other components/layers, e.g., the `AasContributor`, to hook itself into the platform mechanism to create a joint platform AAS.

The *Platform Component* is a collection of the basic services to be started, in particular a (persistent) AAS server or a global network manager. Thus, it requires specific setup information in its `iiecosphere.yml`, e.g., on which port and using which implementation protocol the global platform AAS shall be set up (the individual AAS are then remotely deployed into this AAS server).

Besides the services and their technical network addresses, the platform also uses some pre-defined Transport Layer channels. These channels are briefly summarized in Table 25¹²⁶. It is important that channels are independent of the transport protocol, i.e., apply equally to, e.g., MQTT or AMQP. Moreover, the default metrics channels currently use a fixed JSON format and rely on a default String serializer defined in the Transport Layer. The service channels use an application-specific format determined by the active serializer and the code generation of the platform instantiation process.

Table 25: Transport Channels used by the platform

Channel	Kind	Component	Format	Explanation
EcsMetrics	global	ECS Runtime	JSON	Metrics reporting by ECS Runtime
ServiceMetrics	global	Service Mgt., Service Env.	JSON	Metrics reporting by Services, can be augmented by application-specific metrics.
ComponentStatus	global	<i>all</i>	JSON	Notification about added, changed, removed components. Format pre-defined in transport component.
Trace	global	<i>All</i>	JSON	Optional notification about received/sent data or parameter changes. Enabled

¹²⁵ In Spring applications, this file is typically called `application.yml`. The name for the IIP-Ecosphere platform is different, also as Spring is only used in alternative components.

¹²⁶ Due to problems with Spring Cloud Stream for asynchronous data injection, a workaround for connectors uses further, undocumented streams. In future versions, we plan to completely rely on Spring Cloud Stream to be able to exploit existing binders. If this fails for technical reasons, we will detail the stream names here.

Channel	Kind	Component	Format	Explanation
				per service via configuration model. Format pre-defined in transport component.
Alert	global	<i>all</i>	JSON	Notification about monitoring alerts, format pre-defined in transport component.
<code>data_<service>_<function>_<app>_<instance></code>	global	Application	<i>application specific</i>	Inter-device transport channels per service and function. The channel name also contains the application identifier and, if utilized, the application instance id (may be omitted including the separating underscore).
<code>data_<service>_<function>_<app>_<instance></code>	local	Application	<i>application specific</i>	Intra-device transport channels per service and function (similar to global service transport channels above).

Further, (components of) the IIP-Ecosphere platform recognizes the command line parameters summarized in Table 28.

Table 26: Command line parameters recognized by the platform (state as `--parameterName=value`)

Parameter name	Explanation	Default
<code>iip.port</code>	Network port on which the AAS implementation server shall be running. Overrides the respective value in the platform configuration as well as the environment setting with the same name. May be helpful to setup containers.	<i>Platform configuration</i>
<code>iip.port.svcMgr</code>	Overrides <code>iip.port</code> but only for the service manager.	<i>Platform configuration</i>
<code>iip.id</code>	Device id as used by the ECS-Runtime and the service manager (use consistently, service manager must have same <code>iip.id</code> as corresponding ECS-Runtime). Depending on the strategy realized by the active <code>IdProvider</code> , a device id given on the command line may be ignored. The default provider considers the command line.	MAC-based <code>IdProvider</code>
<code>iip.profile</code>	Start the given lifecycle profile via its name. The default profile starts all lifecycle descriptors of a component. See Table 27 for available lifecycle profiles.	“default”

Table 27: Lifecycle profiles defined by the platform.

Component	Lifecycle Profiles	Default
Monitoring integration for Prometheus	<ul style="list-style-type: none"> Monitoring.prometheus: Start the full integration in an own JVM. Prometheus: Start only the platform-provided Prometheus through an own JVM. Prometheus.exporter: Start only the platform side of the integration, i.e., metrics exporter and alert manger integration in an own JVM. 	Start the full integration with the platform component; conflicting with BaSyx.

Moreover, the IIP-Ecosphere platform recognizes the following more overarching (Java) system properties in Table 28:

Table 28: System properties recognized by the platform (stated as *-Dproperty=value*)

Property	Explanation	Default
<code>iip.pid.dir</code>	Directory where process identification files (PID) in Unix style are stored. These files indicate that a certain process was started. This information can be helpful when automatically starting platform processes, e.g., through a process manager. Default files are <code>iip-platform.pid</code> , <code>iip-ecsRuntime.pid</code> and <code>iip-serviceMgr.pid</code> .	The system's temporary directory.
<code>iip.port</code>	Network port on which the AAS implementation server shall be running. Overrides the respective value in the platform configuration. May be helpful to setup containers.	<i>Platform configuration</i>
<code>iip.port.svcMgr</code>	Overrides <code>iip.port</code> but only for the service manager.	<i>Platform configuration</i>
<code>iip.identityStore</code>	Folder containing the identity store with the file name / format required by the actual pluggable identity store (see Table 24).	.
<code>iip.installedDeps</code>	Folder containing the <code>installedDependencies.yml</code> file required by the service manager and individual services.	.
<code>iip.app.</code>	Name prefix of system properties to be passed on to the running service applications, e.g., to ease debugging.	-
<code>iip.devices.plcNext.grpc.sock</code>	Location of the GRPC socket for the optional PLCnext system-level monitoring plugin.	<code>/run/plcnext/grpc.sock</code>
<code>iip.eclass.keystoreKey</code>	Identity store key providing access to the authentication certificate for the ECLASS web service.	<code>eclassCert</code>
<code>iip.eclass.locale</code>	Locale used as preferred result language when querying the ECLASS web service.	System locale
<code>iip.test.dataFile</code>	Name of the data file in generated service/connector tests. Typically, as documented in the generated code, the default is a service-specific file name, but may be overridden, e.g., for continuous integration.	<i>As documented in generated code</i>

8.3 Compiling the IIP-Ecosphere platform

Due to the various optional and alternative components in the IIP-Ecosphere platform that we manage in individual artifacts/Eclipse projects, compiling the IIP-Ecosphere platform is not trivial. Usually the binaries of the individual components are either available via Maven central (releases) or SSE Maven repository (snapshot, releases). These builds are created and deployed by the SSE Continuous

Integration (CI) server as illustrated in Figure 98, which knows about all the build dependencies among the components and builds the parts and pieces along the dependency tree when the code of a single component changes. As part of building, it executes the respective component tests, assembles the documentation and, if successful, deploys the respective snapshots to the SSE Maven repository or the stable releases from Maven (or related repositories).













✓		IIP_configuration.configuration	17 Minuten - #381	33 Minuten - #379	2 Minuten 46 Sekunden
✓		IIP_connectors	23 Stunden - #216	Unbekannt	47 Sekunden
✓		IIP_connectors_aas	23 Stunden - #231	Unbekannt	44 Sekunden
✓		IIP_connectors_mqttv3	23 Stunden - #233	Unbekannt	57 Sekunden
✓		IIP_connectors_mqttv5	23 Stunden - #237	Unbekannt	52 Sekunden
✓		IIP_connectors_opcuav1	23 Stunden - #218	Unbekannt	55 Sekunden
✓		IIP_deviceMgt	50 Minuten - #30	Unbekannt	34 Sekunden
✓		IIP_ecsRuntime	18 Minuten - #219	Unbekannt	1 Minute 0 Sekunden
✓		IIP_ecsRuntime.docker	14 Minuten - #236	Unbekannt	55 Sekunden
✓		IIP_ecsRuntime.kubernetes	13 Minuten - #120	Unbekannt	41 Sekunden
✓		IIP_monitoring	48 Minuten - #31	Unbekannt	35 Sekunden
✓		IIP_platform	12 Minuten - #282	Unbekannt	37 Sekunden

Figure 98: Screenshot of the SSE Continuous Integration server (IIP-Ecosphere view, cropped)

For completeness, we discuss below the dependencies among the individual components of the IIP-Ecosphere platform (as illustrated in Figure 99). As indicated in Figure 99, some components need specific settings for successful testing, e.g., the RTSA components need to know which JDK to use for RTSA execution (strict Java 8 for the original RTSA).

The `platformDependencies` project collects the dependency information of (optional or required) external components that are used by at least one component and do not constitute singleton wrapped components (cf. Section 4). In other words, the `platformDependencies` project defines the managed dependencies of the platform with their respective version number (range) but without actually using them. The dependent components rely on this information and just state the required components without replicating their version numbers (Maven parent POM mechanism). As we usually do not build external components, e.g., protocols, rather than relying on available release binaries, these dependencies are out of scope.

Building the IIP-Ecosphere platform starts with compiling and packaging the IIP-Ecosphere Maven plugins for dependency management (obtaining, unpacking and updating the configuration meta-model), Maven invocation (executing sequences of goals in different profiles on the same POM) and Python (syntax check and testing). As these plugins are used by the build process defined in the platform dependencies, they must be built before without upstream dependencies.

The Support Components consisting of `support` and the `AAS abstraction support.aas` are the most basic IIP-Ecosphere components without further dependencies to the platform. As `support` provides some basic functionality that is also used in components that strictly depend on Java 8 or would break their build flow, `support` is also based on Java 8. In particular, it integrates a Python helper class identifying the actual Python binary from the IIP-Ecosphere Maven plugin, which, in turn, forces the IIP-Ecosphere Maven Plugin to Java 8.

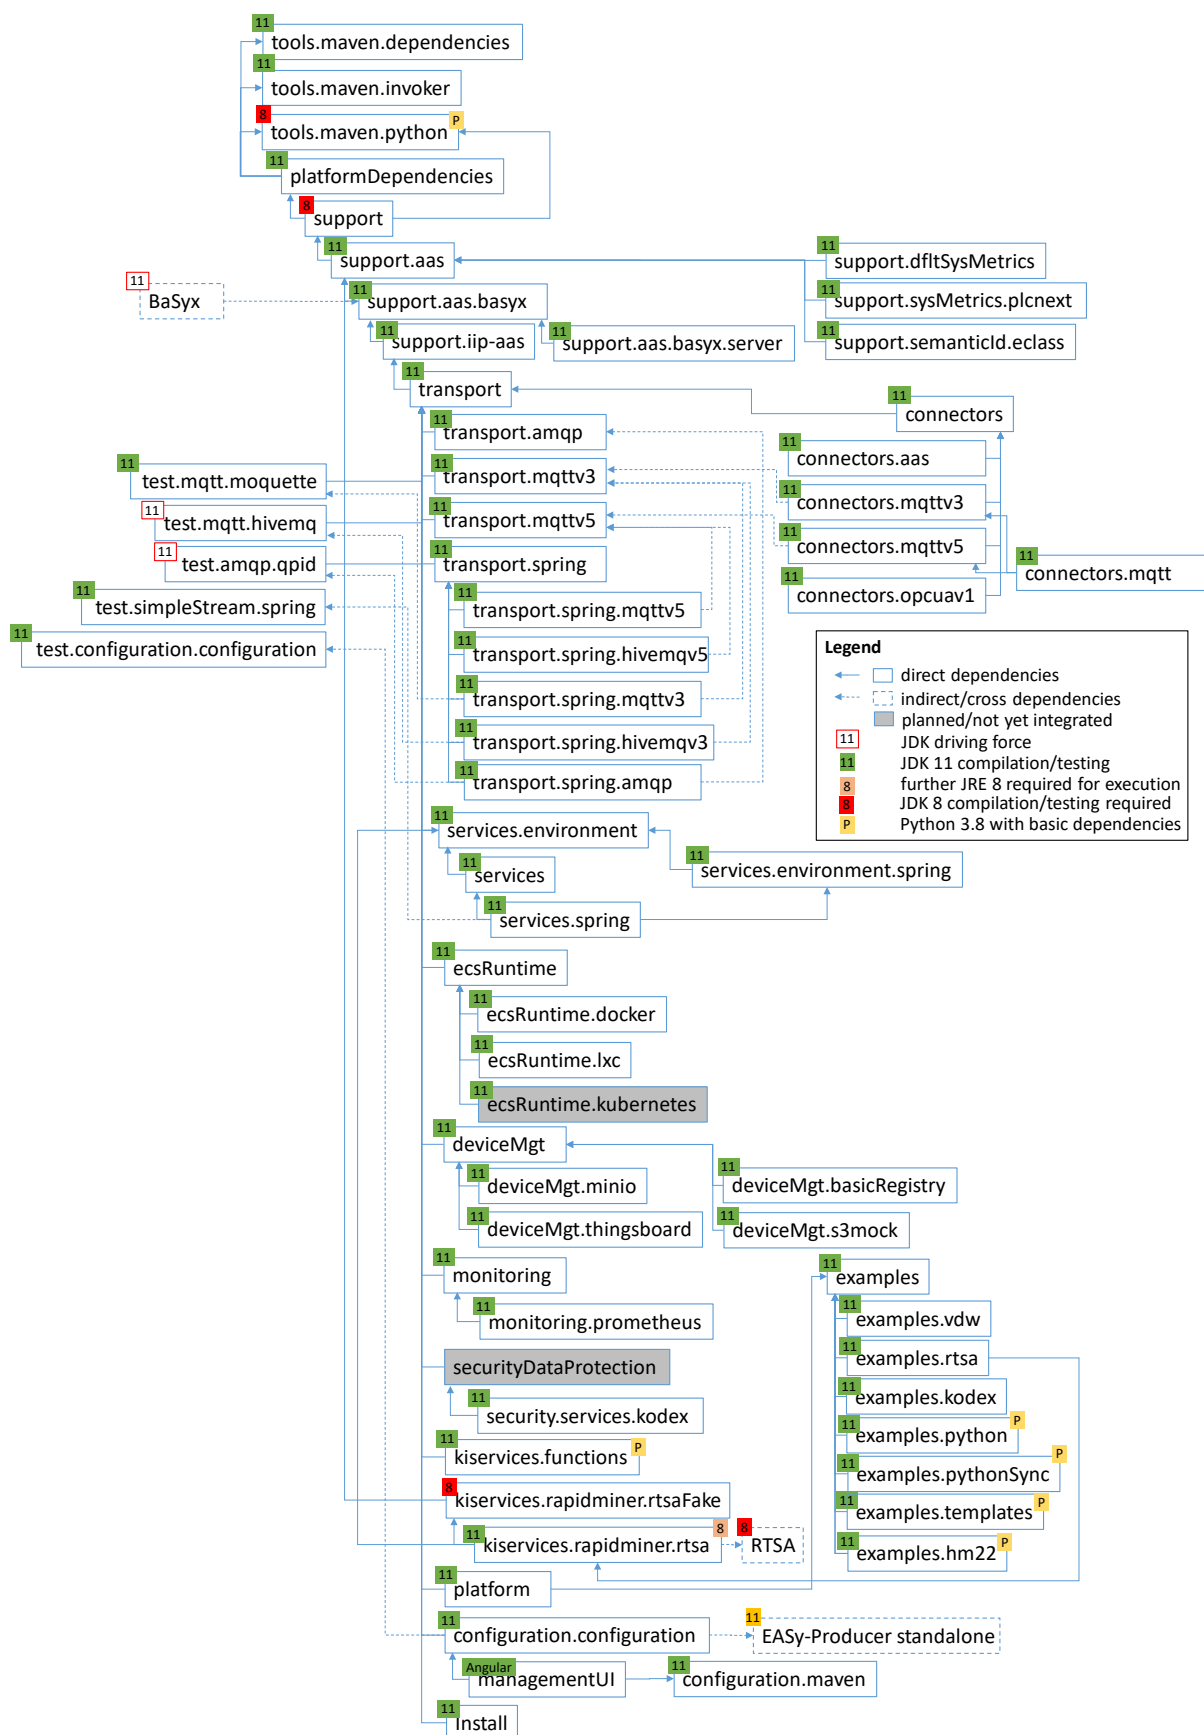


Figure 99: Dependencies among the components (platform examples not shown, folder names in GitHub)

The BaSyx default implementation (BaSyx is considered as main driving force for the selection of the actual Java version of the platform) and the `iip-aas` support functions depend directly on `support.aas` and are build when `support.aas` changes. Further, two basic Maven plugins for realizing an integrated build process for the platform and the applications directly depend (like `support.aas`) on the platform dependencies.

The Transport Component (`transport`) is then the next component to be built after the Support Layer. If `transport` is changed, it triggers the building of the transport connectors (`transport.*`), the basic (optional) Spring integration (`transport.spring`) and the Spring binders (`transport.spring.*`) utilizing the transport connectors. The Connectors Component (`connectors`) relies on the type translation and serialization mechanisms of the Transport Component and, further, the individual platform/machine connectors (`connectors.*`) depend on the Connectors Component. The MQTT platform/machine connectors are, in turn, based on the corresponding transport connectors.

The components of the service layer (`services.*`) consist of the service manager interface including the abstract creation of the AAS (`services`), the specific implementation for Spring Cloud Streams (`services.spring`) as well as the generic service environment (`services.environment`) and the Spring-specific service environment (`services.environment.spring`). The IIP-Maven plugin for Python depends on the generic service environment, in particular to rely on common capabilities to determine the actual Python binary to be executed (which may differ in the CI environment for historical reasons).

The resource/deployment components (`ecsRuntime.*`) are partially realized, e.g., the ECS runtime and the container manager for Docker. The container manager for Kubernetes and the platform monitoring are in planning/realization and not part of this release. The first components for the device management have been integrated (`deviceMgt.*`), including an optional integration of Minio for object storage and ThingsBoard for device management. Similarly, initial steps towards the integration of security and data protection services (`securityDataProtection.*`) have been done. Within the reusable intelligent services, the RapidMiner RTSA (version 14) requires Java 8 for execution. As RTSA is an IPR protected component, the regression testing and the RTSA example are based on a fake version of RTSA, a program that pretends to be RTSA without its AI capabilities. In turn, our fake RTSA must be built with Java 8 to mimic the prerequisites of RTSA.

The platform server(s) component provides the startup sequence for central services as well as the preliminary command line interface for platform functionality. Currently, neither the device management nor the security services are part of the assembled platform server component.

Further, there the integration of the configuration model (`configuration.configuration`), which depends on the capabilities of EASy-Producer (stand-alone, Maven-based integration). The IIP-Ecosphere Maven plugin for executing the platform instantiation depends on `configuration.configuration`.

At the end of the hierarchy, there is currently the IIP-Ecosphere platform management user interface. As the Web user interface is realized in Angular based on information from the platform AAS as backend, this requires a different build process. The TypeScript code of the UI is compiled using angular, packaged, archived by the CI server and then, using a pseudo Maven POM, deployed as binary component into the Maven repository of the platform. The platform instantiation takes this binary up as usual for other components with binary processes, unpacks and customizes the UI.

The Test Components (`test.*`) are a side track but required for testing. The protocol related test components contain integrations of embedded protocol brokers, such as Apache Qpid, HiveMq or

Moquette, which shall be explicit testing dependencies rather than part of the production code. Apache Qpid (due to requested bug fixes) and HiveMq are further driving forces for the selection of the JDK for the IIP-Ecosphere. Moreover, `test.simpleStream.spring` is a testing artifact containing a simple stream processor chain for testing the Spring service manager in the Services Layer. Further, `test.configuration.configuration` provides implementations for the streaming interfaces created by the Test configuration instantiations in `configuration.configuration`. This introduces cyclic dependencies, but only for the very first build. For such a build, the configured artifacts in the respective tests can be set to an empty string, the test passes, creates and deploys the interfaces and the test artifacts can be build and deployed. In a second round, the artifact configuration is restored and the test runs again, now correctly packaging the streaming artifacts.

If a local build is required, we provide a multi-module Maven POM on the top-level of the platform code repository. Given that all required software (Java, Maven, Python and dependencies due to building the examples) is installed as discussed in Section 8.4, the following command (Linux, “~\$” indicates the shell prompt) builds the full platform, the Maven command

```
~$ mvn install
```

builds the full platform. Due to specific dependencies between tests and their service implementation, a first offline build may require the setting `-Diip.build.initial=true` and to speed up the build (otherwise it may take around two hours), tests may disabled with `-DskipTests`. Both optional settings must be stated at the end of `mvn install`.

8.4 Installing and using the IIP-Ecosphere platform

As discussed above, the IIP-Ecosphere platform must be configured and instantiated before it can be executed. Thus, the continuous integration does not provide complete platform dependencies (except for those created as part of testing `configuration.configuration`). Below are the required steps to run the actual release of the IIP-Ecosphere platform **to explore the functionality** (this may change in the future)¹²⁷. The setup was tested with JDK 11, JDK 13, Maven 3.6.3, Maven 3.8.5 and Docker 20.10.7, Windows 10 as well as Ubuntu 20.4.1 and 20.4.3.

1. Prepare the operating system. For the next steps in this section, we assume Linux installed on two machines (assuming 147.172.178.145 as “server” and 147.172.178.143 as “device”, we will adjust the IP addresses in the fourth step). We use the roles “server” (central installation) and “device” (remote compute resource, e.g., Edge device, virtual machine, etc.) to refer to target machines for the installation. In particular, in a test setting, “server” and “device” may be the same physical machine. In the command line illustrations, “~\$” indicates the shell

```
~$ sudo apt install openjdk-13-jdk-headless
~$ sudo apt install maven
~$ sudo apt install docker.io
```

prompt (the output of “\$” depends on your active shell) including the actual directory. We assume that you start in your home directory (~). Install Java JDK¹²⁸ and maven (version 3.6.3). Docker (version 20.10.2) is only needed if you want to execute the platform in a containerized setup.

¹²⁷ Please take care that the containing file path does not contain unusual characters, such as #. This may cause parts of the instantiation to fail.

¹²⁸ A JDK is needed for the platform instantiation rather than a JRE. As also discussed in Section 2, for devices a minimum Java 11 is required, e.g., you may install JDK 13 on the „server“ and JDK 11 on the „device“ or in respective containers. If your installation does not set the `JAVA_HOME` variable, the Maven scripts created during platform installation may issue a stacktrace warning that Javadoc cannot be executed, but the scripts shall pass. To prevent this, set `JAVA_HOME` so that it points to the JDK installation home directory.

On devices, the installation may differ as Java/Maven could be part of the container hosting the ECS runtime or the Service Manager/Services. By default, Docker requires root permissions to execute functions. If you want to use docker as “normal” user¹²⁹, perform a log out and log in so that your group membership is re-evaluated.

```
~$ sudo usermod -aG docker $USER
```

Depending on the use of Python in services, the build process for applications may include Python syntax checking and execution of Python unit tests. If Python is utilized, a recent Python 3.9 with pyflakes and for service execution PyYaml must be installed. Moreover, depending on the utilized platform functions potentially also pyzbar 0.1.9, opencv-python 4.5.5.64, numpy 1.20.1, and Pillow 9.1.0 are required. For the management user interface, Angular 12, express 4.18 and cors 2.8 are required. Below we indicate an example installation of the requirements for Python 3.9 with recommended minimum versions of the dependencies (which may differ on your system).

```
~$ python3 -m pip install pyzbar==0.1.9
~$ python3 -m pip install opencv-python==4.5.5.64
~$ python3 -m pip install numpy==1.20.1
~$ python3 -m pip install Pillow==9.1.0
```

2. On server and devices: Obtain the IIP-Ecosphere platform install package. Snapshots can be obtained from SSE Jenkins for Windows¹³⁰, Linux¹³¹ or from GitHub. The install package for the actual release can be obtained from GitHub. For the further steps, we assume that the install package is stored in the actual directory.
3. On server and devices: Unpack the install package. Watch out that no special characters or whitespaces are in your path to the installation directory as this may cause strange effects,

```
~$ mkdir Install
~$ cd Install
~/Install$ tar xzpvf ../Install.tar.gz
```

particularly on Windows systems.

4. On the server: Edit¹³² the example configuration file TechnicalSetup.ivm1 in src/main/easy so that your local IP address is used. In this release, the devices are not listed in the configuration, i.e., search for 147.172.178.145 and replace this IP by the IP of your server machine¹³³. You may perform more changes, but this requires background knowledge from Section 6 on the platform configuration model. One particular setting to be considered is the globalHost for data transport, the global communication broker. In the example, this shall be automatically set through the global variable platformServer to ease the setup for a single server installation, but in your settings also different servers for different central services (AAS, monitoring, broker, etc.) may be set up, which then require a more individualized setup. Currently, the selection of code artifacts is restricted to the Maven servers used for

¹²⁹ <https://docs.docker.com/engine/install/linux-postinstall/>

¹³⁰ https://jenkins-2.sse.uni-hildesheim.de/view/IIP-Ecosphere/job/IIP_Install/lastSuccessfulBuild/artifact/install.zip

¹³¹ https://jenkins-2.sse.uni-hildesheim.de/view/IIP-Ecosphere/job/IIP_Install/lastSuccessfulBuild/artifact/install.tar.gz

¹³² For the next versions of the platform, we plan to provide UI support for the initial setup as well as the design of applications.

¹³³ For this release we suggest not using HTTPS as schema or a non-empty endpoint path for the AAS server. Also for VAB, HTTPS is currently disabled. We plan to add certificate support in one of the next releases.

development, i.e., further artifacts cannot be obtained from further repositories, e.g., the future platform service store. This will be targeted by one of the next releases.

Depending on the Java version, various settings to open the Java module system may be needed. Currently, these settings are fixed in `IIPEcosphere.ivm1` in the variable `javaOpts` and can be adjusted before generation if needed. Pragmatically, you may also alter the generated shell scripts (cf. steps 5 and 7)¹³⁴.

5. On the server, decide about the broker. You may install one for a certain protocol and adjust the configuration (step 5) accordingly. By default, the configuration goes for an AMQP broker on port 8883. The platform will also instantiate an example broker¹³⁵, i.e., if you want to rely on that broker, there is nothing to do for this step. As alternatives, you can also rely on, e.g., a broker that is shipped with your operating system or with a device.
6. On the server: Instantiate the platform using

```
~/Install$ mvn install
```

If you already instantiated an older (snapshot) version of the platform, please advise maven to update its artifacts, i.e., add the parameter `-U` in run `mvn -U install`. This command executes the `PlatformInstantiator` mentioned in Section 3.11 through Maven, passing it three parameters (implicitly given in the `pom.xml` file), namely the name of the model to instantiate (`PlatformConfiguration`), the relative folder where the model is located (`src/main/easy`) and the folder where to store the instantiated artifacts (the relative folder `gen`). Please note that this may fail if your modifications to the configuration file are syntactically or semantically incorrect. Alternatively, you can check out the full code from GitHub and run the `PlatformInstantiator` from your IDE or force maven to copy all dependencies into a folder and run Java manually on the command line. As several files and folders are produced by the instantiation process, also a `README.txt` file is generated, which provides some explanation on the individual files and folders. Among them are the platform server (`platform`, `plJars`, `platform.*`), the ECS runtime (`ecsRuntime`, `ecsJars`, `ecs.*`) and the service manager (`serviceMgr`, `svcJar`, `serviceMgr.*`). Moreover, to enable the execution on devices with lower resources, the instantiation creates an installation variant containing ECS runtime and service manager running in the same process (`ecsServiceMgr`, `ecsSvcJars`, `ecsServiceManager.*`). In this variant, no container manager is installed (AAS container operations do not exist, executing respective commands through the CLI will lead to an error) and if a container is used as execution environment, the services must run within the same container.

7. On the server: Copy the created artifacts in `gen` (`broker/*`, `ecsJars/*`, `ecs.sh`, `svcJars/*`, `serviceMgr.sh`, `SimpleMeshTestingApp-0.1.0-SNAPSHOT-bin.jar`) to the respective devices (you may leave this step out if server and devices are the same physical machine). For each artifact, the instantiation creates a folder with all dependencies and the respective startup script. In future versions of the platform, this step will be taken over by the device management, the automated container creation and the distribution of containers by the platform.

¹³⁴ Please note that generated files may be overridden when step 6 is executed again.

¹³⁵ For now, the installation package still contains an example AMQP broker, which may be removed in the future. If you want to use that broker, type `cd broker; mvn package; cd ..` to prepare it for execution.

8. On server (and depending on your decision in Step 5): Install and start a protocol broker/server instance. If you decided in Step 6 for the broker instantiated by IIP-Ecosphere, run it in an own shell by

```
~/Install$ cd gen/broker
~/Install/gen/broker$ ./broker.sh
```

9. Start the platform components:
 - On the server: Start the platform server(s) component through the generated platform.sh/bat in the gen folder. If you run all processes from the console as explained in this section, this requires a separate shell for platform, ECS Runtime, service manager and CLI as the processes run intentionally endless (or you start them in parallel/background).
 - Currently, by default, no explicit on/offboarding of devices is needed. If explicit onboarding is required, call the respective ecs or ecsServiceMgr script with parameter onboard.
 - On the devices (if desirable also on the server): Start the broker for local communications, the ECS runtime (ecs.sh/bat) and, finally, the service manager (serviceMgr.sh/bat) through their startup scripts¹³⁶.
10. On server (if all files are available this can also be done on a device): Run cli.sh. An example interactive execution trace was already shown in Section 3.13 in Figure 52. For deploying services, the created application artifact must be on the device running the service manager. Add the artifact to the service manager via the cli (through a local file URL on that device) and start the services. Please note that artifacts and containers are added through their URI, whereby local URIs may differ from system to system¹³⁷, e.g.,
 - Windows: file:///C:/.../SimpleMeshTestingApp-bin.jar
 - Linux: file:/home/ecouser/SimpleMeshTestingApp-bin.jar
 - Service container: file:/apps/SimpleMeshTestingApp-bin.jar
11. On server and devices: To avoid timeouts, a shutdown shall happen in the opposite sequence, i.e., services, service manager, container manager, platform, brokers. Currently, by default, no explicit on/offboarding of devices is needed. If explicit onboarding is required, call the respective ecs or ecsServiceMgr script with the parameter offboard.

If you want to exercise the full cycle, create a Docker container with the service manager and the application artifact first and copy the container to the device running the ECS runtime¹³⁸. It is important to emphasize that these steps shall be automated in future releases.

Copy the container folder from the installation package to your “device”. Copy/move the artifacts from step 7 also into the container folder and execute there

```
~/Install$ docker build -t iip/simplemesh:0.3
-f SimpleMeshTestingApp/Dockerfile .
~/Install$ docker save iip/simplemesh:0.3 | gzip
> SimpleMeshTestingApp/simplemesh-0.3.tar.gz
```

¹³⁶ If you like to take a look into the AAS, open the repository URL

<http://127.0.0.1:9002/registry/api/v1/registry> on the server (please adjust the host/port according to your configuration) and follow the links into the desired AAS submodel.

¹³⁷ The name of the generated service artifacts for Spring changed and is now qualified with the classifier bin. This allows for testing the generated classes prior to any packaging format, which may result in an executable JAR rather than a JAR for testing.

¹³⁸ We provide scripts for creating, saving and running the container as part of the install package.

For convenience, both commands are available as `createAppContainer.sh` and `saveAppContainer.sh` in the install package. At a glance, the second step may appear superfluous, but it is required for the deployment and execution of the container through the ECS runtime. Please take care that the tag `iip/simplimesh` and the file name `simplimesh-0.3.tar.gz` are the same¹³⁹ as in the container descriptor `SimpleMeshTestingApp/image-info.yml`. Add and start the container (similar as described for the services above) through the platform command line interface before starting the services in the container.

With a running platform server and a running ECS runtime, you may also start the container manually. This would then require setting the AAS implementation server port correctly as stated in the container descriptor, i.e., `--network=host --expose port -e "IIP_PORT=port"`. If feasible, you may use the default port 9000 and use `--expose 9000` or more generically `-P` as parameters. An example script is included in the install package as `runAppContainer.sh`.

If you also want to containerize the ECS runtime (one of the possible edge device installations), ensure that the folder `container/EcsRuntime` is on the “device”. For simplicity and to save resources, we map the `SimpleMeshTestingApp` folder as volume into the ECS container (mount point `/SimpleMeshTestingApp`).

```
~/Install$ docker build -t iip/ecsruntime:0.3 -f EcsRuntime/Dockerfile .
~/Install$ docker run -v /var/run/docker.sock:/var/run/docker.sock -P --network=host
--mount type=bind,source="$(pwd)"/SimpleMeshTestingApp,target=/SimpleMeshTestingApp
-it iip/ecsruntime:0.3
```

Akin to the app container, both steps are available as respective scripts in the install package. Before running the ECS container, it is important that the app container has been created and stored. As administrative operations for installing Docker into the container are executed, Docker may issue certain warnings during the creation of the container. The default port for the ECS Runtime AAS implementation server in this Dockerfile is 9000.

For a **permanent installation**, the instantiation process also generates service specifications for Linux systemd, for both, integrated installation of platform services, ECS runtime and service manager on a single machine/container and additional no-dependencies service specifications for device/container installation. The generation of these system service descriptors and their execution is based on several settings in the platform configuration (mentioned below), in particular, as Linux systemd does not take system environment variables into account. These system service descriptors assume an installation of the generated jars (including containing directory, e.g., `plJars`) in the `instDir` (by default `/opt/iip`). For execution, the Java specified in `javaExe` (default `/opt/iip/java`, intended to be a link to the actual JVM) is used. After completion of the startup process, the respective platform executable creates a file containing the process identification (PID) in the `pidDir` (by default `/run` for Linux and the temporary folder for Windows). These PID files are, e.g., considered by the startup scripts of the generated containers.

8.5 Pre-build Docker container images

To ease initial steps with the IIP-Ecosphere platform and to avoid carrying out all the steps described in the last section, we offer two pre-packaged Docker containers on Docker Hub¹⁴⁰,

1. A container with the platform parts (AAS server, the ECS Runtime, the Service Manager) and a simple example service artifact.

¹³⁹ The version number may differ but shall be the same as in the container descriptor.

¹⁴⁰ <https://hub.docker.com/r/iipecosphere/platform>

2. A container with the Command Line Interface.

To experiment with the containers, use the following commands.

Create a network

```
~$ docker network create --subnet=172.19.0.0/16 platform
```

and start the platform parts:

```
~$ docker run --network platform --ip 172.19.0.22 -p 9001:9001
-p 9002:9002 -p 8883:8883 -p 9090:9090 -p 9091:9091 -p
4200:4200 iipecosphere/platform:platform_all.latest
```

Intentionally, the command line interface runs in a separate container, i.e., execute

for the Command Line Interface. An example service artifact is available in the platform container also acting as device under the URI `file:/device/SimpleMeshTestingApp-0.1.0-SNAPSHOT.jar`.

```
~$ docker run -i --network platform iipecosphere/platform:cli.latest
```

The pre-build container allocates the following ports:

- Port `8883` hosts the AMQP broker.
- Port `9001` serves the Platform Asset Administration Shell (AAS).
- Port `9002` provides access to the AAS registry.
- Port `9090` and `9091` hosts monitoring.
- Port `4200` provides UI.

If you want to explore the AAS, open an HTTP browser on port 9002, select the sub-model that you are interested in and click on the link of that submodel, which directs you to an URL on port 9001.

8.6 Considerations for a Distributed Server Installation

As mentioned in Section 8.4, a distributed setup with multiple servers acting in different roles may require a more individualized configuration. In Table 29, we summarize the configuration variables that are of particular importance for such a setup.

Table 29: Summary of configuration variables for a distributed server installation.

IVML Variable	Semantics	Default
<code>transport.globalHost</code>	Global communication broker (depending on transport protocol) for inter-device communication, monitoring, tracing, etc. Default port is 8883, may differ depending on transport protocol.	localhost
<code>transport.localPort</code>	Network port for local broker instances, i.e., brokers for intra-device data transport.	<code>transport.port</code>
<code>transport.localEcsPort</code>	Network port of local broker instances (in ECS-runtime container). Must differ from global transport port if ECS-runtime/service manager shall	8889

IVML Variable	Semantics	Default
	be executed on central servers/broker.	
<code>aasServer.host</code>	Global AAS server host name as used by all devices.	
<code>aasServer.serverHost</code>	Replaces <code>aasServer.host</code> on the server if the AAS server process shall listen to a dedicated network.	-
<code>aasRegistryServer.host</code> <code>aasRegistryServer.serverHost</code>	Akin to <code>aasServer</code> , but for the AAS registry.	
<code>platformMonitoring.host</code>	Server hosting the central platform monitoring, by default Prometheus. May be pre-installed.	127.0.0.1
<code>platformMonitoring.exporterHost</code>	Server hosting the IIP-Ecosphere-to-Prometheus data exporter.	127.0.0.1
<code>platformMonitoring.alertMgrHost</code>	Server running the Prometheus alert Manager. May be pre-installed.	127.0.0.1
<code>managementUI.port</code>	Port of the platform management UI. Currently started on same server as <code>aasServer</code> .	4200

8.7 Environment for Testing and Evaluating the Platform/Applications

With all the features and functionalities provided by the platform and in addition to that, the developed applications and services added to the platform, there is a need to test the platform and the (distributed) applications to make sure that they operate without errors. Moreover, besides validation, it is interesting to evaluate the platform and applications for their runtime properties, e.g., to determine memory consumption, CPU usage, performance and residual capacities for the evolution of the platform as well as individual applications. As the whole cycle of configuring, installing, deploying starting, evaluating, stopping, and uninstalling applications as well as platform components is not trivial, we aim at providing support through the **Platform Evaluation and Testing Environment (PETE)** that allows for automated, distributed, monitored, and repetitive tests and evaluations.

The concept of an experimentation workbench was presented by [38]. They used Jupyter Notebook Project¹⁴¹ which supports experimentation and to some limited degree replication and sharing. Jupyter Notebook in a nutshell is an open-source web-based interactive development environment for notebooks, code, and data. Jupyter Notebook supports over 40 programming languages like Python, R, Julia, and Scala. Using Jupyter Notebook gives the ability to perform a stepwise execution of the experimentation. We decided to use Jupyter Notebook Project to develop and realize the steps in the PETE. In addition to that, there is a script that converts the Jupyter Notebook to a Python script which can be used to do a headless execution (without using Jupyter Notebook).

¹⁴¹ <https://jupyter.org/>

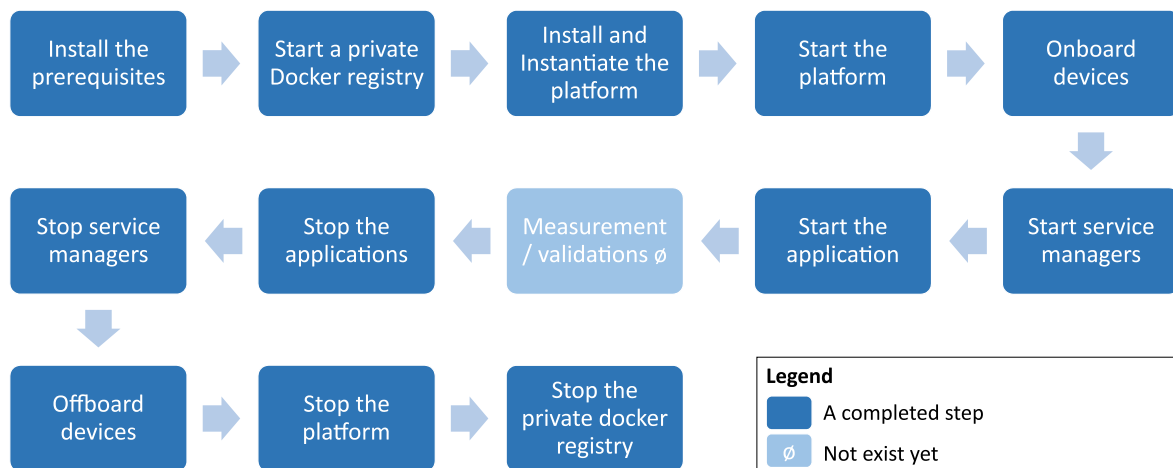


Figure 100: The steps executed automatically by PETE

PETE covers 13 steps as shown in Figure 100: The steps. Those steps are executed automatically by PETE. The step “Measurement / validation” (light-colored step that has “∅” sign) is not implemented yet. The steps are:

1. Install the prerequisites: the required setup needed for the platform on the server and devices that should be used for the test.
 - a. Server setup includes Java, Maven, Docker, and Python with basic dependencies.
 - b. Devices setup; if the device will use the generated containers, then the setup includes just Docker. If the device will use normal scripts, then the setup includes Java, Maven, Docker, and Python with basic dependencies.
2. Starting a private Docker registry: pull and run the registry container, then setup the private Docker registry to not use certificates.
3. Install and Instantiate the platform:
 - a. For the server, it is a full installation of the platform with creating the artifacts and containers for the applications and then adding the containers to the private Docker registry.
 - b. For the device, setup Docker to use the private Docker registry to pull images (specify that no certificate needed for pulling images from the private Docker registry).
4. Start the platform: run the platform scripts on the server (broker.sh, platform.sh, mgtUi.sh, and monitoring.sh).
5. Onboard (join) devices:
 - a. Using the generated containers: pulls the respective ECS container from the private Docker registry and runs it on the device to join the platform.
 - b. The local installation way using scripts: download and run the respective scripts on the device to join the platform (broker.sh, ecs.sh, and serviceMgr.sh).
6. Start service manager: if the device uses the generated containers then run the respective service manager container.
7. Start the application:
 - a. By services functionality in the platform.
 - b. By deployment plan functionality in the platform.
8. Measurement and validations ∅ (not yet implemented).
9. Stop the applications.
10. Stop service managers: if the device uses the generated containers then stop and remove the respective service manager container.
11. Offboard (Remove) devices:
 - a. Using the generated containers: Stop and remove the respective ECS container on the device.
 - b. The local installation way using scripts: Stop the respective scripts on the device (broker.sh, ecs.sh, and serviceMgr.sh).

12. Stop the platform: Stop the platform scripts on the server (broker.sh, platform.sh, mgtUi.sh, and monitoring.sh).
13. Stop the private Docker registry: Stop and remove the registry container.

To use the PETE as stepwise execution using Jupyter Notebook or a headless execution, you should install Jupyter Notebook. The following steps show how to install Jupyter Notebook:

1. If you don't have Python, please install python on your machine.

```
~$ sudo apt install python3.9 -y
```

2. Install Jupyter Notebook from the official website¹⁴².

```
~$ pip install notebook
```

To run Jupyter Notebook and to use the PETE

1. Run Jupyter Notebook, this command will show a URL to use Jupyter Notebook on any browser.

```
~$ jupyter notebook
```

2. Download the PETE from the IIP-Ecosphere GitHub¹⁴³.
3. Copy the PETE to the Jupyter Notebook directory.
4. Open Jupyter Notebook with any browser (using the URL from the step 1) and run the test through.

To do a headless execution of the PETE:

1. Download the PETE from the IIP-Ecosphere GitHub¹⁴³.
2. Copy the PETE to the Jupyter Notebook directory.
3. Execute the following command in the PETE directory to run the test, you might check the logs for the runs in AllLogs folder.

```
~$ ./RunTheTest.sh TestSetup.yaml
```

We are using Python scripts developed in Jupyter Notebook to manage tests. The Python scripts generate command scripts either shell- (Linux) or batch- (Windows) based on the operating system of the machine that runs the Jupyter Notebook, this machine that runs the Jupyter Notebook is called the **TestManager**. The generated command scripts execute the actual scripts remotely to interact with the server and devices. Those scripts are executed remotely to do different actions like installing the prerequisites, starting a private Docker registry etc.

The PETE consists of folders and files:

1. **Setup:** A file (TestSetup.yaml) used to manage the testing environment, the available and usable machines/devices, the platform, the applications to execute, and the tests and their properties (e.g., repetitions, testing/evaluation time, etc.).

¹⁴² <https://jupyter.org/install#jupyter-notebook>

¹⁴³ <https://github.com/iip-ecosphere/platform/tree/main/platform/tests/test.environment>

2. **Scripts:** A folder that holds the shell (Linux) or batch (Windows) scripts that perform the actual execution of the testing/evaluation process on the involved machines. There are two types of scripts - Server and Device scripts - based on the respective role of the machine in the test. Those scripts are moved automatically from the **TestManager** to the machines that are included in the test based on the role of the machine.
3. **Commands:** A folder that holds scripts generated by Jupyter Python Notebook, those commands are generated based on the operating system in the **TestManager**. Those commands remotely execute the scripts moved to the machines.
4. A **Jupyter Python Notebook** (TestManagementScript.ipynb) that reads the TestSetup.yaml setup file then creates and runs the commands which in turn remotely execute the scripts moved to the machines.
5. **Logs:** a folder that holds all the logs from running the commands and scripts.
6. **AllLogs:** a folder that holds the logs from previous runs of the test named by the date and time of the test.
7. **RunTheTest.sh:** a script that run the test in headless execution (Without Jupyter Notebook):
 - a. Convert Jupyter Python Notebook to Python script.
 - b. Extract the number of repetitions of the test from TestSetup.yaml setup file and repeat the test for the number of times defined in the TestSetup.yaml setup file.
 - c. Move the current logs from **Logs** folder to **AllLogs** folder and give it a name (the date and time of the test)

9 How to apply, extend or contribute

In this section, we summarize procedures for some tasks that you may want to perform with the IIP-Ecosphere platform. In the last sub-section (Section 9.5), we provide answers to frequently asked questions.

9.1 Defining an own application-specific service

This is a short/modified form of the explanation in Section 6.7.

1. Adjust your platform configuration¹³³ and define a new service (as discussed in Section 6).
2. Execute the platform instantiation with `generateAppsNoDeps` so that the service interfaces artifact is generated and installed.
3. Import the generated template (`gen/templates`) as project into Eclipse and add only required components as dependencies, in particular your configured application interfaces (see step 1). Alternative and optional components such as AAS implementations or protocols may be added as dependencies in the test scope (see architecture constraints from Section 4).
4. Realize the service, e.g., as Java class(es) in `src/main/java` or Python class(es) in `src/main/python/services` implementing the new interface(s).
 - a. **Java:** For automated creation of instances, services must have two constructors, one taking the `serviceId` and an `InputStream` containing the deployment descriptor YAML as well as a no-argument fallback constructor.
 - b. **Python:** The service must be a class implementing the generated base service, which carries the service identification, description, etc. At the end of your class, the service must instantiate itself to register with the service environment. Further, a Maven artifact descriptor is required to correctly package and deploy the Python sources.

For all supported programming languages, we recommend using the generated application templates, which contain programming frames for Java and Python services, frames for Java and Python tests, testing resources, an example identity store file, required packaging descriptors, and a matching build process (see also Section 6.8).

5. Modify the platform configuration by adding the artifact specification of your service implementation artifact to the configuration of your service(s).
6. Run the platform instantiation with `generateApps` so that the complete artifact is built.
7. Deploy the artifacts to your installation devices, start the platform and try out your service as discussed in Section 8.4.
8. Let IIP-Ecosphere know about your work. In case of a potential open source component, please consider contributing it to IIP-Ecosphere.

9.2 Defining an AAS for a device

Nowadays, AAS is still a rather new approach and devices may not be equipped with their own AAS, although first experimental device plugins do exist, e.g., for Phoenix Contact PLCnext. Currently, as stated in Section 3.8.1, the IIP-Ecosphere allows for two ways of defining an own AAS, in particular to provide a nameplate with device/vendor information or, if desired, a specification of the provided device services.

1. The first approach is to pragmatically state the AAS in a specific Yaml-Format, which is turned into an AAS.
2. The second approach is to provide an AASX package, which makes up the device AAS.

Further forms, e.g., referring to an external AAS, can be realized through plugins. The AAS specification can be:

1. Build into your instantiated platform if the respective files are stated in the resources/device folder before starting the platform instantiation. These files are named according to the device identification, e.g., by default the MAC-Address with all characters in capitals.
2. Loaded from the device upon startup of the respective ECS runtime. Thereby, we search for a file called `nameplate.yml` or `device.aasx` in the installation folder of the platform component, respectively.

The AAS must at least contain a nameplate according to [43, 2].

9.3 Implementing a monitoring/alert data service

Monitoring services are data- or application-specific Java functions that may send out an alert if specific conditions occur. So far, the platform does not implement a generic, e.g., rule-based alert service that can easily be reused. Basically, a monitoring/alert data service is realized as discussed for general application-specific services in Section 9.1. However, to have access to the micrometer monitoring structures, specific steps may be required in Step 4:

Realize the service, e.g., as Java class(es) implementing the new interface(s) and the `MonitoringService` interface from `services.environment`. This will provide a service with access to the `MetricsProvider` in case that you want to add custom metrics. Consider `Transport.sentAlert` if conditions fail. In addition, the interface `UpdatingMonitoringService` is automatically called regularly to update metrics if needed. As above, please consider the rules for constructors.

9.4 Extending the platform by adding a component or a platform service

1. Make yourself familiar with the design of the respective component. Identify the interfaces to implement, e.g., the `Service` interface in `services.environment`.
2. Create a Maven Eclipse project, use the IIP-Ecosphere platform dependencies as parent and add only required components as dependencies. Alternative and optional components such as AAS implementations or protocols may be added as dependencies in the test scope, not in the (default) production scope.
3. Implement your component and test it.
4. Consider extending the platform configuration meta-model, i.e., search for the part describing the components. In some cases, e.g., AAS protocols, this may just be an additional entry in an enumeration. For other components, this may require a new typed IVML compound with default values (akin to the already given compounds). For services, no changes to the meta-model are required.
5. Adjust your platform configuration¹³³ so that your new elements are taken up. In case of a new enum value, use that value. In case of a new compound, replace the existing compound value by a value of your type (providing also the respective settings in the compound value). For a new service, add the service to the application part of your platform configuration and link it into the service mesh (as discussed in Section 6).
6. Run the platform instantiation as discussed in Section 8.4, copy the artifacts to your installation devices, start the platform and try out your extension.
7. Let IIP-Ecosphere know about your work. In case of a potential open source component, please consider contributing it to IIP-Ecosphere.

9.5 Defining a new type in the configuration model

The platform configuration model is equipped with several basic types, e.g., for String, Integer numbers¹⁴⁴, IEC61131-3 date type or OPC UA basic types. However, we do not aim at defining an encompassing set of types rather than mechanisms that allow for adding types later. This is in particular true for record types, which may be highly application specific. This is also true for primitive types. The definition of primitives applies the IVML model pattern for alternatives with detailing properties (Figure 60).

To add a new primitive type, we explain the modifications to the configuration model for the IEC61131-3 date-time type as illustrated in Figure 101:

1. Define a new meta-type for the type, i.e., a compound that just refines the respective base meta type, e.g., `PrimitiveType` or `ArrayType` (lines 3-4 in Figure 101, thus suffix `TypeType`).
2. Define a variable of that type and assign a default (display) name (lines 5-7 in Figure 101). This variable will represent the type, e.g., in record fields and, thus, the variable has the name suffix `Type`. The name will be used by the code generation as fallback alternative in several situations if the behavior is not overridden.
3. Freeze the new variable in order to make it available to the code generation as it is (lines 19-13 in Figure 101).
4. Adjust the intentionally open parts of the code generation. For example, in Figure 101, we extend the language basics for the generation of Java artifacts (lines 16-18) by a specific method which returns the type name to be used (a basic method for `DataType` is defined, the new method is selected through dynamic type dispatch). Similar, the method in lines 21-23 indicates that a type-specific conversion parameter shall be used when obtaining values from a connector parser or writing values to a connector formatter.
5. Adjust also the code generation for Python in `PythonBasics.vt1` (similar to step 4).

¹⁴⁴ Please note that IVML provides similar primitive types that we use to specify the structure of the IIP-Ecosphere configuration model, i.e., the IVML primitive types are on meta (M2) level rather than on model level (M1) that we target here.

```

1 //DataType defined, PrimitiveType refines DataType
2
3 compound IEC61131_3DateTimeTypeType refines PrimitiveType {
4 }
5 IEC61131_3DateTimeTypeType IEC61131_3DateTimeType = {
6     name = "IEC 61131-3 DateTime" // display name, e.g., for UI
7 };
8
9 //...
10 freeze {
11     //...
12     IEC61131_3DateTimeType;
13 };
14
15 //JavaBasics.vtl: toTypeName(DataType) is defined
16 def String toTypeName(IEC61131_3DateTimeTypeType type) {
17     "java.util.Date"
18 }
19
20 //DataOperationBasics.vtl: convName(DataType) is defined
21 def String convName(IEC61131_3DateTimeTypeType type) {
22     "Date" // operation name infix for toDate and fromDate
23 }
24
25 //DataOperationBasics.vtl: convParam(DataType) is defined
26 def String convParam(IEC61131_3DateTimeTypeType type) {
27     "'DT#'YYYY-MM-dd-HH:mm:ss.SS"
28 }

```

Figure 101: Adding IEC 61131-3 date time to the primitive types of the configuration model

Figure 101 may create the impression that all foundational parts of the model must be modified to introduce new data types. While this is true for default types that shall be shipped with the platform, it is not necessarily true for data types needed by individual applications. Thanks to the dynamic dispatch across imported IVML/VTL modules, it is possible to define the type and the overridden functions also in their own platform configuration module, which imports the shipped platform model.

However, it is important that also certain consistency rules apply. The types declared in `JavaBasics.vtl` and the names/conversion parameters declared in `DataOperationBasics.vtl` must match the underlying implementation, in particular of the `InputConverter` and `OutputFormatter` interfaces declared in the `Connectors Component`. There, the specified Java types must be valid input/output parameters of operations with name infix declared in the corresponding `convName` function. For the used Java types, the assumption is that types declared in the Java library or in the imported IIP-Ecosphere components can be used. Other types must be converted by the connector or service targeted by the code generation. Moreover, if a conversion parameter is required, e.g., for date or time values, also the value provided by the `convParam` function must match. For date/time conversion parameters, formatter strings as valid for the Java `SimpleDateFormat` class or logical names as defined in the `FormatCache` class of the connectors component may be used, e.g., `ISO8601` for a specific date format not supported by `SimpleDateFormat`. Further additions to `JavaType.vtl` may be required if the type is declared as primitive type and realized in Java as object. Then the approaches to obtain a hashcode and an equals comparison may have to be adjusted from Java primitive types to object types.

9.6 Using a different transport protocol

The default transport protocol that is being set up and integrated by the platform integration is AMQP. The following describes what steps are needed to use one of the alternatives, e.g., MQTT v5.

1. Open your configuration for editing, e.g., `src/main/easy/InstallTest.ivml` in the install package explained in Section 8.4.
2. Search for the variable `transportProtocol` and replace in the instance type the suffix AMQP by MQTTv5 (or MQTTv3, respectively) as illustrated in Figure 102. AMQP requires an authentication which is obtained from the identity store using the key “amqp”. In contrast, MQTT does not require an authentication (which may differ depending on your setup). The values for `globalHost`, the global communication broker (by default “localhost” or in `InstallTest.ivml` explicitly set to `globalhost`) may be adjusted as needed.
3. Add the `serviceProtocol` assignment (last line of the lower box in Figure 102), which is so far set to AMQP by default in the shipped models.
4. Re-run the instantiation of your platform to enact the changes.

```
transportProtocol = TransportProtocolAMQP {
    globalHost = serverHost,
    port = 8883,
    security = {
        authenticationKey = "amqp"
    }
};

transportProtocol = TransportProtocolMQTTv5 {
    globalHost = serverHost,
    port = 8883
    // no security by default, depends on setup
};

serviceProtocol = ServiceProtocolMQTTv5 {};
```

Figure 102: Switching the transport protocol from AMQP to MQTT.

The platform may instantiate an example broker for the selected protocol using `generateBroker` on as argument for the `PlatformInstantiator`. Please consider that different combinations of brokers and protocol implementations may imply different timing behaviour, e.g., when an operation shall be considered as failed due to timeouts. For this purpose, the transport protocol defines the `actionTimeout` setting, a timeout in milliseconds when an operation is considered as failed. By default and experience, a sufficient timeout for AMQP is 1000 ms and for MQTT 3000 ms (considered automatically as minimum values during platform/application instantiation).

9.7 Observe or debug the data processing

To figure out what is happening in the data processing in a service mesh, of course you can add respective logging calls to the code that implements your services. Please consider that depending on the approach such additional calls may impact the performance (throughput), i.e., approaches to at least disable the logging (as usually supported by logging frameworks) shall be applied. In addition, the platform currently offers two further approaches that are generated into the service glue code that is calling your code and, thus, also applies to connectors and services that are not realized by your organization: Logging and tracing.

- *Logging* emits information via the default logging framework used in the platform as debug information to the respective log output (as set up the console, a file, etc.).
- *Command line streams* (sysout) logging is like logging via a framework, but emits the information without filtering (and specific logging formatting). This may be helpful if it is

inconvenient to change the logging level of the logging framework or other software takes control over filtering of the logs and prevents debug output.

- *Tracing* means sending trace entries through the trace channel (see Section 8.1) to any interested party, in particular the hybrid `TraceToAasService` (see Section 3.7.3.1). The `TraceToAasService` may act as a sink in a service mesh and turns in particular received trace entries into an AAS structure. That service can be used as it is (generic service) or customized by code to allow, e.g., for application-specific operations.

All services in the configuration model offer individual tracing settings as declared in the Services module of the configuration model as illustrated in Figure 103.

```
enum TraceKind {NONE, TRACE, LOG, SYSOUT};

abstract compound ServiceBase /* ... */ {
    //...
    TraceKind traceRcv = TraceKind::NONE;
    TraceKind traceSent = TraceKind::NONE;
    //...
};
```

Figure 103: Switching the transport protocol from AMQP to MQTT.

Currently, there are four modes, no tracing (NONE) as well as tracing (TRACE) and logging as debug information via the logging framework (LOG) and logging on the system output stream (SYSOUT). These modes can be set for data reception (`traceRcv`) and data sending (`traceSent`) of a service. If enabled, respective code is inserted during code generation. By default, tracing is disabled on all services, but you may activate it individually, e.g., to focus debugging on a certain part of a service mesh.

In some cases, knowledge on the input / output of services may be sufficient to identify a problem. In some cases, the problem is part of the routing of the data among the services. For this purpose, the application allows to set a debug flag (`Boolean debug = false`), which may enable additional logging and debugging methods of the service execution framework, e.g., for Spring Cloud Stream.

Enabling debugging requires a re-instantiation of the application. Usually, the application rather than its interfaces is sufficient, i.e., `generateApps` rather than `generateAppsNoDeps` would be the required option for the `PlatformInstantiator` (see also Section 6.2).

9.8 Frequently Asked Questions (FAQ)

In this section, we summarize some questions and issues that repeatedly occurred.

9.8.1 Error parsing HTTP header

Symptom: A part of the platform (platform server, ECS runtime, service manager or platform command line interface) issues an exception with the following message:

```
org.apache.coyote.http11.AbstractHttp11Processor.process Error parsing HTTP
request header Note: further occurrences of HTTP header parsing errors will
be logged at DEBUG level.
```

Reason: One reason may be that a client such as the command line interface tries to access a platform server (AAS server, registry) with an encrypted protocol (HTTPS) while the server is running a non-encrypted protocol (HTTP).

Solution: Ensure that the certificates for client and server side do match. For this release, do not run the platform with an encrypting protocol¹³³.

9.8.2 Maven artifact missing

Symptom: While working with the platform against a release version in Maven, it appears that one of the (non-java) artifacts is missing.

Reason: Although we carefully check the artifacts before a release, it may be the case that the automatic deployment (script) missed some.

Solution: Please let us know about the problem via the IIP-Ecosphere website or via GitHub.

9.8.3 XXX has been compiled by a more recent version of the Java Runtime

Symptom: While executing in particular central parts of the platform, this error message/exception may occur.

Reason: Maven tends to resolve dependencies to the most recent version using a given version number as minimum, in particular if version ranges are allowed. As long as dependencies do not change or the specified version range is feasible, no such problems shall occur. It may occur upon the first resolution, i.e., during installation or when dependencies are updated, e.g., during CI when Maven is requested to search for more recent snapshots. However, in particular for Eclipse components which declare version ranges, compiler settings currently seem to be changed from JDK 8 to JDK 11, i.e., even a minor version change may suddenly (upon an unintended update) lead to this failure.

Solution: As quick fix, use the JDK version indicated by the class version number to execute the respective platform component, i.e., usually JDK 11. Alternatively, you may force Maven to download a compatible, previous version number by creating a simple POM and deleting the failing version. In any case, please let us know about the problem via the IIP-Ecosphere website or via GitHub. We will try to fix the version numbers for central parts such as Eclipse components to a version range that allows for safe execution.

9.8.4 Platform code cannot be setup in Eclipse, e.g., parent POM missing

Symptom: Your IDE reports missing Maven artifacts and shows compilation errors, in particular the parent POM of the platform is missing. Similarly, the code style checking may fail due to missing style definition file.

Reason: The parent POM of the platform defines the versions of non-singleton/wrapped libraries (cf. Section 4). Without that particular POM, compilation cannot run successfully as the artifact version numbers/ranges are missing. If you are working with a release version, it may also be the case that one of the released artifacts is missing (cf. Section 9.8.2).

Solution: Please refer to the code setup guide in GitHub¹⁴⁵.

9.8.5 Unknown platform coding conventions

Symptom: After a first contact with the platform code it seems that you are missing detailed information about applied conventions on how to write code and you cannot find all conventions in this document.

Reason: Although we tried to capture the most important conventions in this document, this document is not intended to be a programmer's guide, i.e., we do not necessarily repeat all coding conventions here.

Solution: Please refer to the platform coding guidelines in GitHub¹¹⁸.

¹⁴⁵ <https://github.com/iip-ecosphere/platform/blob/main/platform/documentation/Guideline.pdf?raw=true>

9.8.6 Maven does not find app dependencies

Symptom: When executing the platform instantiation, Maven complains about missing dependencies.

Reason: Typically, Maven dependencies for apps that ship as examples with the platform are deployed into a public Maven repository. If Maven is executed locally, sometimes required artifacts are not deployed correctly.

Solution: Please open a shell, navigate into the respective directory of the app or the app installations and run `mvn install`.

9.8.7 Execution of application fails due to `Java CompileError`

Symptom: An app is built correctly but when starting it, a `Java CompileError` occurs and prevents the app from starting up. The messages indicate that packages are missing that are actually in the app fat jar.

Reason: We observed this, if app implementation projects override compile settings defined by the platform dependencies. In more details, when you create a new app implementation project, your IDE initially does not know that you will rely on the IIP-Ecosphere platform dependencies and may sets up the compile settings according to your local IDE compile settings. When inheriting your POM from the platform dependencies to rely on IIP-Ecosphere build setup, the IDE compile settings remain in your POM and in extreme cases may conflict with class files included from the generation. This can lead to a `Java CompileError` (a kind of class loading link error).

Solution: Please remove any local compiler setup from your POM files and run Maven on app implementation and app project again.

9.8.8 Services do not start due to problems with `javax.el.ExpressionFactory`

Symptom: When starting services, the service manager reports class loading or instantiation problems for `javax.el.ExpressionFactory`.

Reason: The Java Expression Language (EL) is required by Hibernate, which in turn is used by Spring Boot/Cloud Stream. The EL ships in two parts, interfaces and implementation. Over the time, several versions and implementations of both parts occurred. The Spring packages used by the platform declare a dependency to `jakarta.el`, which ships both parts in the same jar (version 3.0.3). However, due to transitive dependencies, e.g., to Tomcat from BaSyx, further versions such as `tomcat-el-api` or the original `javax.el` may be parts of service implementation dependencies. Multiple versions of the interfaces may, dependent on the classloading sequence, interfere and cause the described symptom. In the generated parts, we try to prevent such overlaps, but, however, we cannot be aware of the dependencies declared by your implementing classes.

Solution: Identify all interfaces and implementations of EL and exclude superfluous ones from the dependencies. Alternatively, try to enforce a class loading sequence that loads `jakarta.el` before all other EL interfaces and implementations. Typically, the generated parts and the default Spring service packaging take care of that. Similarly, the ZIP service artifact including an explicit classpath file are packaged to consider this issue, but due the use of wildcards for ZIP service artifacts not containing a classpath file, an intended class loading sequence cannot be guaranteed then.

9.8.9 Service execution through platform fails

Symptom: When starting services through the platform, the service manager reports a state change to FAILED instead of RUNNING.

Reason: There are multiple reasons that can cause this symptom:

1. Failures in the service code, e.g., a Python-based service cannot be started because the required Python script is not correctly packaged or cannot be executed due to implementation errors.
2. Communication failures as the network communication is not set up correctly, e.g., required ports are already used or not accessible (firewall, not declared as external ports in containers, container is not running in host network mode, etc.)
3. Timing issues in particular when services are started (the first time) in a container.

Solution: Depending on the actual reason, e.g., failures in service code must be solved or communication failures can be addressed by correct network configuration (including the respective settings in IIP-Ecosphere container descriptors). Timing issues often occur when the waitingTime for the service manager is not set correctly. The default value is 1 minute, but on resource-constrained devices, 2 or 3 minutes may be more adequate.

9.8.10 Why do platform scripts always check for recent dependency snapshots

Symptom: When starting platform services/applications or executing build commands, the Maven build process always checks for most recent snapshots. This is time consuming and leads to long console output.

Reason: As long as the platform is in development, it is convenient for us to have the most recent snapshot builds available. This is in particular true for the continuous integration. Always checking for the most recent snapshots is only enabled on the IIP-Ecosphere Maven snapshot repository, not for other Maven repositories.

Solution: You can disable this behavior allowing Maven, e.g., to check for snapshots only once a day. Go to your local Maven repository (usually in your home directory in the folder .m2) and modify the settings file there. If there is no settings file, you can create a new one as shown below.

```
<settings xmlns=http://maven.apache.org/SETTINGS/1.1.0
  xmlns:xsi=http://www.w3.org/2001/XMLSchema-instance
  xsi:schemaLocation="http://maven.apache.org/SETTINGS/1.1.0
    http://maven.apache.org/xsd/settings-1.1.0.xsd">

  <repositories>
    <repository>
      <id>SSE</id>
      <name>SSE</name>
      <url>https://projects.sse.uni-hildesheim.de/qm/maven/</url>
      <layout>default</layout>
      <snapshots>
        <updatePolicy>daily</updatePolicy>
      </snapshots>
    </repository>
  </repositories>

</settings>
```

If there is already a repositories section, please add the contents for the “SSE” repository as shown above. The important setting is the update policy, set this to daily (by default, from the platform dependency, the value is always).

10 Summary & Conclusions

Realizing an open (experimental) IIoT/I4.0 platform is a significant amount of work. IIP-Ecosphere is performing that work and this whitepaper provides technical insights into the ideas, concepts, rationales, designs and implementation state of the current release of the IIP-Ecosphere platform. The rationale behind this document is to enable interested parties to discuss with IIP-Ecosphere on a technical level, to try out the platform or to provide extensions. As the platform is evolving, this document is just a snapshot in time. Future versions may include lessons learned and (reactions to) feedback in order to improve the platform and also this document.

We discussed the technical basis for architecture modeling, the overview of the layered architecture, the individual layers and the components they contain. For each component, we provided a requirements analysis [based on [13, 41]] and a discussion of the realized requirements. We discussed architectural constraints, the actual use of Asset Administration Shells (AAS), the approach to platform configuration and instantiation, future contributions to the (external) security of the platform, selected implementation details as well as how-to's on applying and extending the platform.

To conclude about the actual state of the realization, we provide below some insights into selected realization Key Performance Indicators (KPI), namely requirements fulfillment, connectors, developed components, testing, use of open source components, and use of Asset Administration Shells.

Table 30 summarizes the discussed and realized requirements. The platform handbook of the current version discussed more than half of the top-level and sub-requirements for the platform. As several components are not yet realized or in realization but not part of this release, i.e., for which we do not discuss the requirements (status), we can also conclude that about a third of the requirements are already either completely or partially realized (and tested).

Table 30: KPI-based summary of discussed/realized requirements

KPI: Requirements (from [13], 141 top-level and 181 sub-requirements)	
Discussed top-level requirements	107 (76% of all top-level requirements)
Discussed sub-requirements	126 (70% of all sub-requirements)
Completely realized requirements	119 (37% of all requirements)
Partially realized requirements	24 (10% of all requirements)

From the requirements reported above, 43% of all requirements have been realized. Regarding the open requirements, to mention the categories with highest numbers, 33 (10% of all) are on data storages/semantic data integration, 30 (9% of all) on data protection, 25 (8% of all) on the AI toolkit, 21 (7% of all) on virtualization/containers, 17 (5% of all) on applications and application management, 13 (4% of all) on security and 11 (3% of all) on runtime adaptation. However, it is important to mention that also important functionality demanded by practical experience with the platform has been realized, that is just mentioned as a few/single requirement(s), e.g., the platform management UI, generated implementation templates, generated service tests, or the container instantiation have not been mentioned in the requirements, e.g., the distributed testing and evaluation environment (PETE), the example or demonstration applications.

Table 31 summarizes the number of “connectors” realized so far. In particular, various basic protocol connectors for transport and streaming have been realized and tested. For some components, even more connectors do exist, e.g., for Spring Cloud Stream.

Table 31: KPI-based summary of realized connectors

KPI: Connectors (*requiring application-specific extensions)	
Support Layer AAS connector	3 including the default implementation for BaSyx, 1 device measurement “connector”, 1 semantic ID resolver

KPI: Connectors (*requiring application-specific extensions)	
Spring transport connectors	8 binders: RabbitMQ, Kafka, Kafka Streams, Amazon Kinesis, Google PubSub, Solace PubSub+, Azure Event Hubs, Apache RocketMQ
IIP-Ecosphere transport connectors*	3 connectors: MQTT v3, MQTT v5, AMQP 5 connector binders: MQTT v3 (Paho, Hive), MQTT v5 (Paho, Hive), AMQP
Machine/platform connectors*	4 connectors: OPC UA v1, AAS, MQTT v3, MQTT v5
Storage connectors	1 connector to Amazon S3 (local, remote; two realizations)
Sum	16 connectors realized, 8 further available
Security	transport connector extensions for IDS
Data integration	generic database connector
Cloud connectors	semantic-based optional cloud connectors (if within resources, AWS and Gaia-X)
Application	northbound external platform connectors for data exploration and linking of IIP-Ecosphere platform instances
Further planned	6

Table 32 summarizes the number of developed components categorized by the layers or logical components.

Table 32: KPI-based summary of developed components

KPI: Components developed	
Tooling	4 (maven plugins)
Support Layer	8 (including 2 optional component)
Transport Component	10 (including 8 optional/alternative components)
Connectors Component	5 (including 4 optional/alternative components)
Services Layer	4 (including 2 optional/alternative components)
Resources Layer	11 (including 7 alternative components)
Configuration Layer	1
Platform server(s) Component	1
Platform management UI	1
Sum	45 (including 25 optional/alternative)

Table 33 summarizes the number of test cases realized by the platform. For judging the overall number, it is important to recall that the granularity of tests differs significantly, ranging from classical unit tests over integration tests up to validation and instantiation of a configuration in a single test. Also, the number of tests differs, e.g., in the Services Layer, many fine-grained monitoring tests from [3] are defined that increase the number significantly. Moreover, the number of test cases is only one side of the testing medal. It is also import to consider coverage metrics. The line coverage is typically between 69% and 89% except for the following: The Spring environment is currently not directly tested rather than indirectly via the Spring service testing artifact, the test components that either consist of testing code only or define an artifact for component testing, e.g., the service testing artifact, or that are currently not part of the release (Kubernetes resource manager, device management, platform monitoring).

Table 33: KPI-based summary of tests

KPI: Tests (of various granularity, from unit to integration)	
Support Layer	178
Transport Component	49
Connectors Component	30
Services Layer	163
Resources Component	147
Security and Data Sharing Component	2
Reusable Intelligent Services Component	5
Configuration Component	20 + 4 (generated)
Platform server(s) Component	5
Tooling, Installation support	9
Examples	3 + 7 shell-based build and run tests
Sum	765 test cases (including 4 generated cases), usually 48%-89% line coverage (exceptions for specific components with acceptable reasons)

Table 34 summarizes the number of open source components used in and integrated into the platform. It is important to mention that Table 34 just lists top-level components and not their transitive dependencies. For example, the platform server component ultimately consists of more than 200 libraries that stem from the transitive dependencies of 9 top level component dependencies and one auxiliary dependency.

Table 34: KPI-based summary of used open source components

KPI: Used Open Source Components (only distinct/top-level ones are listed)	
Support / AAS factory connector	2 BaSyx, as alternatives Jsensors
Transport component	4 Eclipse Paho, HiveMQ client, Rabbit MQ client, Spring Cloud Stream 4 for testing: Apache Qpid Broker J, HiveMq, Googlecode JSON simple, Google protobuf, Moquette
Connectors component	2 Apache Milo (and as above Apache Paho) Testing relies on the components mentioned above as well as the server implementations provided by the used components.
Services component	2 Micrometer (with Spring Cloud Stream), Spring Cloud Stream Local Deployer
Resources component	10 Docker, Java Docker client, ThingsBoard, MinIO, S3Mock, Prometheus, Prometheus Java client, Prometheus alertmonitor, LXC, JLXD
Security and Data Protection	1 KODEX
Reusable intelligent services	3 Rapidminer RTSA, zxing, pyzbar
Configuration	1 EASy-Producer
SUM	23 in production code, 4 for testing
Planned: Data Lakes	At least one feasible database
Planned: Security	IDS/GAIA-X
Planned: Installation	Broker like Eclipse Mosquitto or RabbitMQ
Planned: AI	Flower
Further planned	at least 7

We provide a coherent Asset Administration Shell for the platform on each installed device, e.g., through the ECS runtime installations. Thus, the number of individual AAS, which differ due to the heterogeneity of the devices, depends on the actual on-site installation and leads to **$r+s+a+1$ linked AAS** with r being the number of devices with ECS runtime installations, s the number of services, a the number of applications and 1 central platform AAS in a factory (assuming a remotely deployed AAS). From a type perspective, this leads to two AAS types, one for the ECS installations and one for the central IT installation. Instead of accounting for that number, we count the number of sub-models (more precisely **sub-model types**) contributing to the IIP-Ecosphere platform AAS.

Table 35: KPI-based summary of realized asset administration shells

KPI: Asset Administration Shells and sub-models	
Support Layer	1 dynamic sub-model (types)
Transport Component	1 static sub-model (transport connectors)
Connectors Component	2 dynamic sub-models (installed connectors, active connectors), while the active connectors change dynamically at runtime
Services Component	3 active sub-models (services, artifacts, relations), change dynamically at runtime, 1 AAS per service with at least 1 sub-model
Resources Component	3 active sub-models (resources, containers, device management), 1 AAS per device with at least 1 sub-model
Configuration Component	0 not realized so far
Applications	1 AAS with 3 sub-models
Platform (via Support Layer)	1 sub-model, 1 platform AAS, 1 status/progress AAS
SUM	12 sub-models in platform production code, 1 AAS for the platform, 1 AAS/sub-model per service, 1 AAS/sub-model per device, 1 AAS with 3 sub-models per application
Further planned	at least one per layer

One interesting KPI is the number of available services of the platform. We do not count here for platform-integrated basic services such as the service manager, the ECS runtime or the platform servers rather than for services that can be integrated into applications. As discussed in this handbook, we must distinguish generic (platform-provided and application-provided) as well as application-specific services. Table 36 summarizes the current state. Although the number of overall services may appear to be rather low, it is important to recall that these services are generic and automatically adjusted to the utilizing service mesh. Moreover, the RapidMiner RTSA service itself is generic and can execute AI services defined in RapidMiner as well as service chains defined in RapidMiner studio. While, the number of instantiable services is unlimited, we are aware of the fact that further services shall be integrated in the future, e.g., from the IIP-Ecosphere demonstrators or the AI-accelerator.

Table 36: KPI-based summary of the provided services

KPI: Implemented (AI) services	
Connectors	3 platform-provided, generic connector types with model-based integration for OPC UA, MQTT v3 and v5. Instantiated connectors are integrated as services into service meshes.
Security	1 platform-provided, generic anonymization and pseudonymization service (KODEX)
AI	1 platform-provided, generic AI service (RapidMiner RTSA)
Applications	1 platform-provided, hybrid service acting as application endpoint for mirroring application specific information into

KPI: Implemented (AI) services	
	the application AAS. In application-specific form, this service may act as command endpoint.
HM'22/TddT'22 demonstrator	2 generated connectors, one re-used hybrid application AAS service. Further, three application-specific services including a QR-code detecting camera source, a Python AI service and an action decider.
SUM	9 in production code, > 10 application-specific services for testing
Further planned	> 10

In summary, the second basis release accompanied by this platform handbook already realizes roughly a third of the intended functionality and, thus, provides a good basis for platform research and case studies. However, also basic functionality that would be desirable for certain work is still missing. Thus, for the next release, we plan in particular for the following missing functionality:

- Improved UI including access to configuration
- Optional integration of Kubernetes based on flexible protocols
- Integration of further (AI) services.

11 References

- [1] A. S. Ahmadian, Model-based privacy by design, PhD thesis, University of Koblenz and Landau, Germany, 2020.
- [2] S. Bader, H. Bedenbecker, M. Billmann, A. Bondza, B. Boss, S. Erler, K. Garrels, T. Hadlich, M. Hankel, O. Hillermeier, M. Hoffmeister, M. Kiele-Dunsche, J. Neidig, A. Orselzki, S. Pollmeier, B. Rauscher, W. Rieder, S. Stein, B. Waser, Generic Frame for Technical Data for Industrial Equipment in Manufacturing (Version 1.1), Plattform Industrie 4.0, 2020, https://www.zvei.org/fileadmin/user_upload/Presse_und_Medien/Publikationen/2020/Dezember/Submodel_Templates_of_the_Asset_Administration_Shell/201117_I40_ZVEI_SG2_Submodel_Spec_ZVEI_Technical_Data_Version_1_1.pdf
- [3] M. G. Casado, Service and device monitoring on devices in IIP-Ecosphere, IT-Studienprojekt, Universität Hildesheim, 2021
- [4] M. G. Casado, H. Eichelberger, Industry 4.0 Resource Monitoring - Experiences with Micrometer and Asset Administration Shells, Symposium on Software Performance 2021, EUR Workshop proceedings
- [5] J.-H. Cepok, Projektarbeit, Uni Hildesheim, 2023 (in preparation)
- [6] J.-H. Choi, J. Park, H. D. Park, O. Min, DART: Fast and Efficient Distributed Stream Processing Framework for Internet of Things, ETRI Journal, 39 (2), pp. 202-211, 2017
- [7] R. Dautov, S. Distefano, D. Bruneo, F. Longo, G. Merlino, A. Puliafito, Pushing Intelligence to the Edge with a Stream Processing Architecture, International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData), pp. 792-799, 2017
- [8] H. Eichelberger, A Matter of the Mix: Integration of Compile and Runtime Variability, *Workshop on Dynamic Software Product Lines, FAS'16*, 2016.
- [9] H. Eichelberger, C. Qin, K. Schmid, Experiences with the Model-based Generation of Big Data Applications, Lecture Notes in Informatics (LNI) - Datenbanksysteme für Business, Technologie und Web (BTW '17), S. 49-56, 2017
- [10] H. Eichelberger, S. El-Sharkawy, C. Kröher, K. Schmid, IVML Language specification, http://projects.sse.uni-hildesheim.de/easy/docs-git/docRelease/ivml_spec.pdf
- [11] H. Eichelberger, G. Palmer, S. Reimer, T. Trong Vu, H. Do, S. Laridi, A. Weber, C. Niederée, T. Hildebrandt, Developing an AI-enabled IIoT platform - Lessons learned from early use case validation, SASI4'22 @ ECSA'22, 2022
- [12] H. Eichelberger, G. Palmer, C. Niederée, Developing an AI-enabled Industry 4.0 platform - Performance experiences on deploying AI onto an industrial edge device, Symposium on Software Performance (SSP'22), 2022
- [13] H. Eichelberger, C. Sauer, A. S. Ahmadian, M. Schicktanz, A. Dewes, G. Palmer, C. Niederée, IIP-Ecosphere Platform – Requirements (Functional and Quality View), Version 1.0, March 2021, IIP-2021/02-en, DOI: 10.5281/zenodo.4485774, 2021
- [14] H. Eichelberger, C. Qin, R. Sizonenko, K. Schmid, Using IVML to Model the Topology of Big Data Processing Pipelines In Proceedings of the International Systems and Software Product Line Conference SPLC'16, p. 204 – 208, 2016.
- [15] H. Eichelberger, K. Schmid, EASy Variability Instantiation Language: Language Specification, http://projects.sse.uni-hildesheim.de/easy/docs-git/docRelease/vil_spec.pdf

- [16] X. Fu, T. Ghaffar, J. C. Davis, D. Lee, EDGEWISE: A Better Stream Processing Engine for the Edge, USENIX Annual Technical Conference, pp. 929-945, 2019
- [17] E. Gamma, R. Helm, R. Johnson, J. Vlissides, Design Patterns: Elements of Reusable Object-Oriented Software, Addison-Wesley, 1995
- [18] Ø. Haugen, Common Variability Language (CVL) – OMG Revised Submission, OMG document ad/2012-08-05, 2012
- [19] C. Hochreiner, M. Vögler, P. Waibel, S. Dustdar, VISP: An Ecosystem for Elastic Data Stream Processing for the Internet of Things, pp. 19-29, EDOC'16, 2016
- [20] C. Hochreiner, M. Vögler, S. Schulte, S. Dustdar, Elastic Stream Processing for the Internet of Things, CLOUD, 2016
- [21] J.-H. Hoepman, Privacy design strategies - (Extended Abstract). In ICT Systems Security and Privacy Protection - IFIP TC 11 International Conference (SEC'14), pages 446–459, 2014.
- [22] International Data Spaces, IDS reference architecture model version 3.0, <https://internationaldataspaces.org/22m-3-0/>
- [23] The Industrial Internet Reference Architecture Technical Report, <https://www.iiconsortium.org/pdf/IIRA-v1.9.pdf>
- [24] J. Jürjens, Secure Systems Development with UML, Springer, 2005
- [25] H. Koziolok, S. Grüner, J. Rückert, A Comparison of MQTT Brokers for Distributed IoT Edge Computing, ECSA, 2020
- [26] LNI 4.0 Testbed Edge Configuration – Usage View, https://www.plattform-i40.de/PI40/Redaktion/EN/Downloads/Publikation/LNI4.0-Testbed-Edge-Configuration_UsageViewEN.pdf
- [27] E. Maleki, F. Belkadi, N. Boli, J. van der B. Zwaag, K. Alexopoulos, S. Koukas, M. Marin-Perianu, A. Bernard, D. Mourtzis, Ontology-Based Framework Enabling Smart Product-Service Systems: Application of Sensing Systems for Machine Health Monitoring, IEEE Internet of Things Journal, 5 (6), pp. 4496-4505, 2018
- [28] N. Martz, J. Warren, Big Data - Principles and best practices of scalable realtime data systems, Manning, 2015
- [29] D. O’Keeffe, T. Salonidis, P. Pietzuch, Frontier: Resilient Edge Processing for the Internet of Things, VLDB Endowment, 11 (10), pp. 1178-1191, 2018
- [30] OMG, Unified Modeling Language, Version 2.5.1, <https://www.omg.org/spec/UML/About-UML/>
- [31] D. Pidun, Geräteverwaltung von IoT-Geräten für die IIP-Ecosphere Plattform, BSc-Abschlussarbeit, Universität Hildesheim, 2021
- [32] Plattform Industrie 4.0, Die Verwaltungsschale im Detail, 2019, <https://www.plattform-i40.de/PI40/Redaktion/DE/Downloads/Publikation/verwaltungsschale-im-detail-pr%C3%A4sentation.html>
- [33] Reference Architecture Model Industrie 4.0, <https://www.plattform-i40.de/PI40/Redaktion/EN/Downloads/Publikation/rami40-an-introduction.html>
- [34] B. Satzger, W. Hummer, P. Leitner, S. Dustdar, Esc: Towards an Elastic Stream Computing Platform for the Cloud, IEEE International Conference on Cloud Computing, pp. 348-355, 2011
- [35] C. Sauer, H. Eichelberger, A. Ahmadian, A. Dewes, J. Jürjens, Aktuelle Industrie 4.0 Plattformen – Eine Übersicht, IIP-Ecosphere Whitepaper IIP-2020/001, 2020, DOI: 10.5281/zenodo.4485756, 2020

- [36] K. Schmid, H. Eichelberger, EASy-Producer: From Product Lines to Variability-rich Software Ecosystems, SPLC' 15, 2015
- [37] K. Schmid, S. El-Sharkawy, C. Kröher, Improving Software Engineering Research through Experimentation Workbenches. arXiv e-prints, arXiv-2110, 2021
- [38] L. Schulz, Container-Virtualisierung mit LXC in der IIP-Ecosphere-Plattform, BSc Arbeit, Uni Hildesheim, 2023
- [39] M. Staciwa, Experimentelles Container-Deployment auf Industrie 4.0 Geräte, Projektarbeit, Uni Hildesheim, 2020
- [40] M. Staciwa, Modell-basierte Erstellung von containervirtualisierter Industrie 4.0 Anwendungen am Beispiel der IIP-Ecosphere-Plattform, Bachelorarbeit, Uni Hildesheim, 2022
- [41] H. Stichweh, C. Sauer, H. Eichelberger, IIP-Ecosphere Platform Requirements (Usage View), Version 1.0, Januar 2021, IIP-2021/001, DOI: 10.5281/zenodo.4485801, 2021
- [42] F. van der Linden, K. Schmid, E. Rommes, Software Product Lines in Action - The Best Industrial Practice in Product Line Engineering, Springer, 2007
- [43] ZVEI, Specification Submodel Templates of the Asset Administration Shell – ZVEI Digital Nameplate for industrial equipment (Version 1.0), https://www.plattform-i40.de/IP/Redaktion/DE/Downloads/Publikation/Submodel_Templates-Asset_Administration_Shell-digital_nameplate.html
- [44] T. Ziadi, L. Hélouët, J.-M. Jézéquel, Towards a UML profile for Software Product Lines, Intl. Workshop on Software Product-Family Engineering, 2003

Über die Autoren



Dr. Holger Eichelberger is deputy head of the Software Systems Engineering group at the Institute of Computer Science at the University of Hildesheim. He conducts research in the areas of software product lines, model-based engineering, performance monitoring, and performance analysis. In particular, he is interested in the integration of these areas to create adaptive software systems. In IIP-Ecosphere he leads the think tank "Platforms" as well as the AI Accelerator. He studied computer science at the University of Würzburg, where he received his PhD on the automatic layout of UML diagrams.

Fotograf: Daniel Kunzfeld



Dr. Amir Shayan Ahmadian is a postdoctoral researcher at the Faculty of Computer Science at the University of Koblenz-Landau. His research interests focus on the challenges of designing and implementing secure and privacy-friendly software systems as well as on the current developments in Industry 4.0. He studied computer science at the University of Paderborn and received his PhD in computer science from the University of Koblenz-Landau. During his doctorate, he developed a methodology to operationalize the principle of "data protection through technology design".



Dr. Andreas Dewes holds a PhD in experimental quantum computing from the Sorbonne University of Paris and the French Nuclear Energy Agency (CEA). He has founded several software companies and is the CEO of KIProtect GmbH, which develops advanced technical software solutions for data protection and data security. Within IIP-Ecosphere, KIProtect GmbH is developing a solution for the secure and privacy-compliant use of industrial & IoT data together with the consortium project partners and associated companies.



Marco Ehl is a research associate in the Software Engineering working group under the direction of Prof. Dr. Jan Jürjens at the Institute for Software Technology at the University of Koblenz-Landau. He researches model-driven methods for software development. His focus is on the analysis and explainability of automated production systems. He obtained his Master of Science degree in computer science at the University of Koblenz-Landau on the topic of model-based monitoring of integrated state machines.



Ahmad Alamoush is a research associate in the Software Systems Engineering group at the Institute of Computer Science at the University of Hildesheim. He conducts research on deployment of containerized applications and the adaptation of the deployment at runtime. He obtained his Master of Science degree in computer science at the University of Jordan on the topic of the 0-1 knapsack problem on chained-cubic tree interconnection network.



Monika Staciwa studies computer science at the University of Hildesheim. In IIP-Ecosphere, Monika works in particular on container management, (virtualized) asset administration shells, the python service environment and (automatic) creation of containers.



Miguel Gómez Casado studies computer science at the university of Valladolid. During his ERASMUS+ visit at the University of Hildesheim, Miguel worked on service monitoring, representing monitoring information in and querying monitoring information from asset administration shells.