EMERALD

Deliverable D2.9

Runtime Evidence Extractor – v2

Editor(s):	Angelika Schneider, Florian Wendland
Responsible Partner:	Fraunhofer AISEC (FHG)
Status-Version:	Final v1.1
Date:	15.09.2026
Type:	OTHER
Distribution level:	PU

Project Number:	101120688
Project Title:	EMERALD

Title of Deliverable:	D2.9 – Runtime Evidence Extractors – v2
Due Date of Delivery to the EC	31.10.2025
Updated Version of Deliverable	15.09.2026

Workpackage responsible for the Deliverable:	WP2 – Methodology for knowledge extraction
Editor(s):	Angelika Schneider, Florian Wendland (FHG)
Contributor(s):	-
Reviewer(s):	Verena Geist (SCCH) Cristina Martínez, Juncal Alonso (TECNALIA)
Approved by:	All Partners
Recommended/mandatory readers:	WP1, WP2, WP3, WP4, and WP5

Abstract:	This deliverable presents tools and techniques for evidence extraction from runtime information that can be integrated with the Certification graph. It is the result of work performed in Task 2.5. This document is the final version on runtime evidence extractors, highlighting all changes to the first/interim version reported in D2.8.
Keyword List:	Evidence collection, runtime information, cloud, *Clouditor-Discovery*, *Codyze-Provenance,* technical evidence.
Licensing information:	This work is licensed under Creative Commons Attribution-ShareAlike 4.0 International (**CC BY-SA 4.0 DEED** https://creativecommons.org/licenses/by-sa/4.0/)
Disclaimer	Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union. The European Union cannot be held responsible for them.

Document Description

Version	Date	Modifications Introduced	
		Modification Reason	Modified by
v0.1	06.10.2025	First draft version, ToC	Angelika Schneider (FHG)
v0.2	06.10.2025	Add content to sections 1.3, 1.4, 3	Angelika Schneider (FHG)
v0.3	06.10.2025	Add content to section 4	Florian Wendland (FHG)
v0.4	07.10.2025	Finalizing document for the internal review	Angelika Schneider (FHG)
v0.5	17.10.2025	Internal QA review	Verena Geist (SCCH)
v0.6	28.10.2025	Incorporate QA review comments	Angelika Schneider, Florian Wendland (FHG)
v0.7	29.10.2025	Final review	Cristina Martínez, Juncal Alonso (TECNALIA)
v0.8	30.10.2025	Incorporate final review comments	Florian Wendland, Angelika Schneider (FHG)
v1.0	31.10.2025	Submitted to the European Commission	Cristina Martínez, Juncal Alonso (TECNALIA)
v1.0.1	15.06.2026	Include updates of all sections, e.g., • Status update of requirements for *Clouditor-Discovery* and *Codyze-Provenance* • Updates to the limitations of *Clouditor-Discovery* and *Codyze-Provenance* • Inclusion of *Clouditor-Discovery* for IONOS • Addition of a summary figure for *Clouditor-Discovery* • Clarification on the context and scope of *Codyze-Provenance*	Angelika Schneider, Florian Wendland (FHG)
v1.0.2	25.06.2026	Internal QA review	Verena Geist (SCCH)
v1.0.3	08.07.2026	Incorporate QA review comments	Angelika Schneider, Florian Wendland (FHG)
v1.1	15.09.2026	Submitted to the European Commission	Cristina Martínez, Juncal Alonso (TECNALIA)

Table of contents

List of tables

List of figures

Terms and abbreviations

AI	Artificial Intelligence
AI-SEC	AI Security Evidence Collector
AMOE	Assessment and Management of Organisational Evidence
API	Application Programming Interface
AWS	Amazon Web Services
CertGraph	Certification Graph
CI/CD	Continuous Integration / Continuous Deployment
CLI	command-line interface
Codyze	Static Code Analyzer from FHG
CSAF	Common Security Advisory Framework
CSP	Cloud Service Provider
eknows-e3	*eknows* evidence extractor
GA	Grant Agreement to the project
GitLab	Version control and DevOps platform
gRPC	Google Remote Procedure Call
in-toto	Framework defining attestation format for software supply chains
KPI	Key Performance Indicator
KR	Key Result
MEDINA	Predecessor project of EMERALD
Protobuf	Protocol Buffers
REST	Representational State Transfer
SLSA	Supply-chain Levels for Software Artifacts
TRL	Technology Readiness Level
WP	Work Package

Executive Summary

This deliverable presents the final design, architecture, and implementation state of the runtime evidence extractors *Clouditor-Discovery* and *Codyze-Provenance* of WP2. It contributes to the key result KR1-EXTRACT of EMERALD, a framework to continuously extract information from different layers of a cloud service and prepare suitable evidence based on them.

EMERALD follows a knowledge graph-based approach to provide a unified view of the cloud service under certification at different layers of the service, ranging from the infrastructure layer (e.g., virtual resources) to the business layer (e.g., policies and procedures), to the implementation layer (e.g., source code files) and data layer (e.g., increasingly used AI models) in cloud applications. The runtime evidence extractors, developed in Task 2.5 *Extraction of evidence using runtime information* and described in this deliverable, aim on the one hand at identifying critical security-related functionality such as data encryption, transport encryption, or authentication in cloud infrastructure components. On the other hand, they establish software provenance and artefact attestation to completely track software from its inception as source code to deployed build artefacts. This is complementary to the evidence gathered in Task 2.2 *Extraction of evidence / knowledge out of source code*. Other related deliverables in WP2 provide functional and technical details on further evidence extractors from different sources, i.e., D2.3 [1] on source evidence extraction in Task 2.2, D2.5 [2] on evidence extraction from policy documents in Task 2.3, and D2.7 [3] on security and privacy preserving evidence extraction in Task 2.4. All these details contributed to D2.10 [4] and D2.11 [5] on the overall information model of the certification graph in Task 2.1.

This document starts by illustrating how the runtime extractors fit into the overall EMERALD architecture. The main part provides functional and technical descriptions of the evidence extractors *Clouditor-Discovery* and *Codyze-Provenance*, including their purpose and scope, the (current and planned) coverage of the EMERALD requirements, and the components' internal architectures. These descriptions are complemented by information on delivery and usage, as well as on limitations. Finally, the document concludes with a brief summary.

The runtime evidence extractors described in this deliverable contribute to KR1-EXTRACT by providing next-generation evidence gathering tools and techniques based on a knowledge graph approach. The extractor *Clouditor-Discovery* is implemented and integrated with other components of the EMERALD architecture. The requirement of the component is already satisfied by the presented prototype. The second extractor *Codyze-Provenance* is a new component within the EMERALD architecture and is also integrated in the EMERALD architecture.

This is the second iteration of the deliverable coming from Task 2.5. Evidence is prepared according to the integrated, graph-based model of semantically linked and combined evidence, provided in D2.10 (interim version) [4] in project month 15 (January 2025) and D2.11 (January 2026) [5] in project month 27 (January 2026). The extracted evidence is stored and assessed, i.e., to verify the implementation of security metrics, in the scope of WP3 [6].

1 Introduction

EMERALD aims to provide a next generation set of evidence gathering tools and techniques based on a knowledge graph approach. KR1-EXTRACT supports an improved and unified tool-supported approach to continuously extract knowledge from different layers of a cloud service, e.g., infrastructure, platform, runtime information, policy documents, software, and AI models.

The objective of WP2 is to establish a unified view of the cloud service being certified, known as the target of evaluation, by extracting and enriching knowledge of the different layers of the service and providing suitable evidence for security metrics. A major part of this work package is research and design of multiple tools and techniques to extract knowledge out of various sources. A graph-based model, called the certification graph (*CertGraph*), serves as a common structure that is filled by all evidence extraction components.

1.1 About this deliverable

The goal of this deliverable is to present the design and implementation of the EMERALD evidence extractors that extract runtime information from cloud services. This is a report on the prototype reflecting the implementation and integration of the extractors and is the second of two iterations of deliverables, resulting from Task 2.5 *Extraction of evidence using runtime information*, with the first deliverable being D2.8 [7].

Evidence on the runtime information level is gathered by the runtime information evidence extractor *Clouditor-Discovery*, which supports the *CertGraph* data model. The *Clouditor-Discovery* component is based on the respective microservice of *Clouditor*[1] and was already used in MEDINA[2]. It focuses on generating evidence for security-related findings, such as encryption in use, at rest encryption or restricted ports. While the component was at TRL 5 in MEDINA, it should be advanced to a higher TRL in EMERALD.

Another source of runtime information are CI/CD pipelines and their jobs. During the build of a cloud service or application from source code, jobs such as application security testing, software composition analysis, and secure software development measures augment the confidence in the security of the final build artefact. *Codyze-Provenance* is a new addition to EMERALD as part of *Codyze* that gathers evidence about CI/CD pipeline executions. This evidence facilitates to assess security enhancing jobs executed during a build in a CI/CD pipeline. Moreover, *Codyze-Provenance* assists attestations for executed jobs and link them together to provide provenance. From attestation and provenance reports it becomes possible to track the complete supply chain of software artefacts from source code to deployed artefacts.

Furthermore, application-specific runtime information (e.g., found in log files) might be used to provide additional context regarding the executed functionality. However, we have decided not to pursue this approach, as the log files do not yield valuable information beyond what we already receive from other queries.

1.2 Document structure

The document is structured as follows.

In Section 2, we discuss how the runtime evidence extractors fit into the overall EMERALD architecture and their relationship with other components. Section 3 describes the *Clouditor-Discovery* evidence extractor, which provides the extraction of runtime information of cloud services. Section 4 describes the *Codyze-Provenance* evidence extractor, which provides

[1] https://github.com/clouditor/clouditor/tree/main/service/discovery
[2] https://medina-project.eu/

traceability and attestation along CI/CD pipelines. For each extractor, we provide functional and technical descriptions, along with information about delivery, usage, and limitations.

Section 5 ends up with the conclusions of this deliverable.

1.3 Updates from D2.8

The deliverable is based on D2.8 [7] and presents an updated version. It includes sections from D2.8 that remain unchanged, along with newly added information. To facilitate tracking the updates and enhancements made since the previous version, Table 1 details the modifications and new additions for each section of the document. Please also note that the latest version v1.1 of this deliverable reflects the final status of development of the runtime evidence extractors until the end of WP2.

Table 1. Overview of the changes to D2.8

Section	Changes
1 Introduction	Added new subsections: • "Updates from D2.8" and • "Technological Advances from the MEDINA project"
2 Runtime evidence extractors in the EMERALD architecture	Minor updates, text moved to new section 1.4.
3 *Clouditor-Discovery*	• Added updates to the new evidence extractor for OpenStack and IONOS. • Updated requirement CLDISC.01 • Added updates to the evidence extractor for IONOS • Updated limitations and future work
4 *Codyze-Provenance*	• Moved *Codyze-Provenance* from design phase into early prototype for testing and evaluation. • Updated all sections to reflect new status of the component. • Clarified use of SLSA provenance and in-toto attestations • Updated requirement(s) • Added required user support for SLSA and in-toto • Updated limitations and future work
5 Conclusion	Minor updates.

1.4 Technological Advances from the MEDINA project

EMERALD builds upon the outcomes of the MEDINA[3] project: *A security framework to achieve a continuous audit-based certification in compliance with the EU-wide cloud security certification scheme* (GA 952633). While MEDINA only reached Technology Readiness Level (TRL) 5, the goal of EMERALD is to enhance the system and elevate its TRL from prototype (TRL5) to product (TRL7). This will be achieved by

- making AI and high-level technologies approachable for non-experts,

[3] https://cordis.europa.eu/project/id/952633

- increasing the abstraction level in evidence extraction across various layers of the cloud service (infrastructure, code, policy documents) in an easy-to-use and understandable way, and
- providing trustworthy support for continuous auditing processes.

To facilitate this, we develop an interaction concept that connects various components through a new UI/UX for the overall EMERALD product and ensures alignment with the auditors' workflow.

In MEDINA, some evidence extraction components implemented their own assessments, leading to increased maintenance efforts when requirements changed over time. Therefore, centralizing assessment is one of the goals in EMERALD. This is achieved by delivering exclusively (or as far as possible) raw evidence to the *Evidence Store* [6]. These changes and extensions affect the components listed in Table 2, which shows an overview of the tools from both projects along with the extensions made during the transition.

Table 2 Changes in components from MEDINA to EMERALD

EMERALD tool	MEDINA tool	Changes
Clouditor-Discovery	Discovery (part of the *Clouditor*)	• Added new evidence extractors for the OpenStack and the IONOS Cloud. • Send evidence to the *Evidence Store* rather than to the *Assessment* component. • Usage of the owl2proto tool [8].
Codyze-Provenance	None, newly developed in EMERALD	• Collect provenance reports from CI/CD pipelines. • Send evidence to the *Evidence Store* and *Trustworthiness System* regarding software provenance.

2 Runtime evidence extractors in the EMERALD architecture

This section describes how the runtime evidence extractors interact with (selected) EMERALD components on a conceptual level. Figure 1 shows the EMERALD high-level architecture as a component diagram. In EMERALD, a component is defined as "any part of the EMERALD ecosystem that has a specific functionality and can be considered a separate entity with respect to other components" [9]. In contrast, a tool is a "software element that has several disparate functions and therefore can be composed by several components" [9]. Therefore, *Clouditor-Discovery* and *Codyze-Provenance*, as a part of *Codyze* [1], are referred to as components rather than tools in the context of EMERALD.

The components for collecting evidence about technical and organisational measures, i.e., *AMOE*, *eknows-e3*, *AI-SEC, Clouditor-Discovery*, and *Codyze*, are represented at the bottom part of Figure 1. The source evidence extractor components *Codyze* and *eknows-e3* [1] obtain technical evidence from the analysis of the source code of cloud applications. *AMOE* [2], the component for organisational evidence gathering from MEDINA, analyses various documents and policies of the cloud service provider (CSP) and produces evidence about the CSP's compliance to organisational requirements of the certification framework. *Clouditor-Discovery*, also originated from MEDINA, collects evidence about the secure configuration of cloud resources, with a focus on runtime data extraction in EMERALD. *AI-SEC* [3] is a newly developed component in EMERALD and analyses AI models for several key evidence regarding robustness against adversarial attacks, explainability, and fairness. *Codyze-Provenance* is a new addition to *Codyze* to support runtime evidence extraction, therefore in the EMERALD component diagram, it is part of the represented *Codyze* component.

As previously mentioned in Section 1.4, evidence is now sent to the *Evidence Store* instead of the *Assessment* component. All WP2 extraction components, which extract knowledge from the various layers of a cloud service (i.e., policy documents, source code, cloud interfaces, AI models, etc.), provide (part of) evidence (e.g., for transport encryption), which is then mapped to the EMERALD evidence format using the terms described in the *CertGraph Ontology* [5]. This evidence information is stored in the *Evidence Store* following the defined schema and is used to assess the metrics defined in the *Repository of Controls and Metrics* [6].

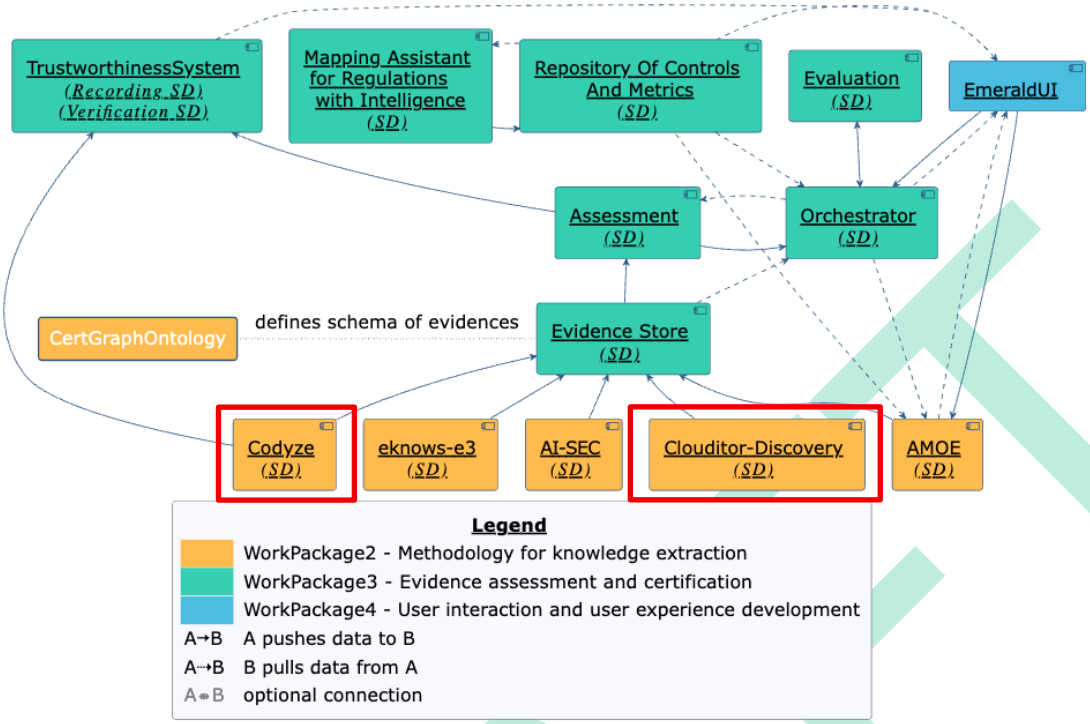

Figure 1. Excerpt of the EMERALD component diagram [10]. The highlighted component Clouditor-Discovery and the component Codyze-Provenance, which is part of the Codyze component, are described in this deliverable.

3 Clouditor-Discovery

The *Clouditor-Discovery* component is responsible for discovering security-relevant configurations from cloud resources. As such, it is one of the evidence extractors that can be integrated into the EMERALD framework. The following sections describe the component in detail regarding the overall purpose, important requirements, architecture, technical aspects, usage and the limitations.

3.1 Functional description

Overall purpose. The component *Clouditor-Discovery* serves the purpose of identifying and discovering security-related configurations across various cloud resources within the EMERALD project. Its primary goal is to enhance security compliance by discovering security-relevant configurations, such as encryption in use, encryption at rest and restricted ports.

Context and scope. *Clouditor-Discovery* is a microservice within the *Clouditor*[4] tool, which was successfully used in the MEDINA project[5] [11]. In EMERALD, the *Clouditor-Discovery* component differs from MEDINA in the variety of discovered resources (e.g., Azure Key Vaults, Functions, Web Apps) and the use of the *Owl2proto*[6] tool (see D2.1 [12]). Additionally, it has been further developed to support additional discoverers, such as OpenStack[7]. Other components of the *Clouditor* used in EMERALD, such as the *Orchestrator*, *Assessment* and *Evidence Store*, are described in Deliverable D3.4 [6].

The *Clouditor-Discovery* focuses on cloud resources, including Virtual Machines, Object Storage, and Network Interfaces, from multiple CSPs like Azure. It achieves this by utilizing API calls to collect runtime information, which is mapped to the EMERALD evidence format in accordance with the *CertGraph* Ontology [5], and then stored in the *Evidence Store* component. With the newly developed tool *Owl2proto*, as outlined in D2.1 [12] and the related paper [8], we can seamlessly integrate the EMERALD evidence format through automatically generated proto files from the Ontology directly into the *Clouditor-Discovery* component. This approach enhances efficiency and significantly reduces maintenance efforts compared to the previous manual process for creating and updating Ontology objects.

Motivation. A key goal of the EMERALD project is to assess the cloud service security, making evidence from these services crucial for the assessment. To achieve a comprehensive result across all layers of the cloud services, evidence extractor components are necessary to analyse as many layers as possible. The *Clouditor-Discovery* serves as the key component for extracting runtime configurations from cloud services, building upon its integration in the earlier MEDINA architecture. It has been enhanced to include new discoverers for additional cloud providers, further expanding its capabilities. In addition to this component, other evidence extraction components for various other layers of the cloud services have been integrated in EMERALD, including source code, organisational and AI model evidence extractors.

[4] https://github.com/clouditor/clouditor
[5] https://medina-project.eu/
[6] https://github.com/oxisto/owl2proto
[7] https://www.openstack.org/

Requirements. The technical requirement for the *Clouditor-Discovery* component with its implementation progress (degree of implementation - *partially / fully /not implemented*) is provided in Table 3.

The *Clouditor-Discovery* component is currently implemented for the CSPs Azure, AWS, Kubernetes, OpenStack[8], and IONOS[9]. However, the implementation for OpenNebula[10] is no longer planned, as evidence from OpenNebula is now sent directly to the *Evidence Store*.

Table 3. Requirement CLDISC.01 - Discovery of security features of infrastructure components

Field	Description
Requirement ID	CLDISC.01
Short title	Discovery of security features of infrastructure components
Description	The *Clouditor-Discovery* needs to discover security properties of infrastructure components. The evidence with the security properties is sent to the *Evidence Store* in the ontology format.
Status	Validated
Priority	Must
Component	*Clouditor-Discovery*
Source	KPI
Type	Technical
Related KR	KR1_EXTRACT
Related KPI	KPI 1.1
Validation acceptance criteria	Validated if evidence arrives at the *Evidence Store*.
Progress	100% - Fully implemented
Milestone	MS6: Integrated audit suite V2 (M30)

3.2 Technical description

In this section, we give the technical description of the runtime evidence extractor component *Clouditor-Discovery*.

3.2.1 Prototype architecture

The *Clouditor-Discovery* is a component of the *Clouditor* tool which employs a microservice architecture allowing individual components to scale or to add new components, e.g., adding several evidence collection components for new cloud services. Figure 2 shows an overview of the *Clouditor-Discovery* component. The component has several available evidence extractors for Cloud Providers, generates evidence based on the gathered runtime information and

[8] https://www.openstack.org/
[9] https://cloud.ionos.com
[10] https://opennebula.io/

submits the evidence to the *Evidence Store* for further processing. The components are written in Go[11] and communicate among each other via the gRPC[12] protocol.

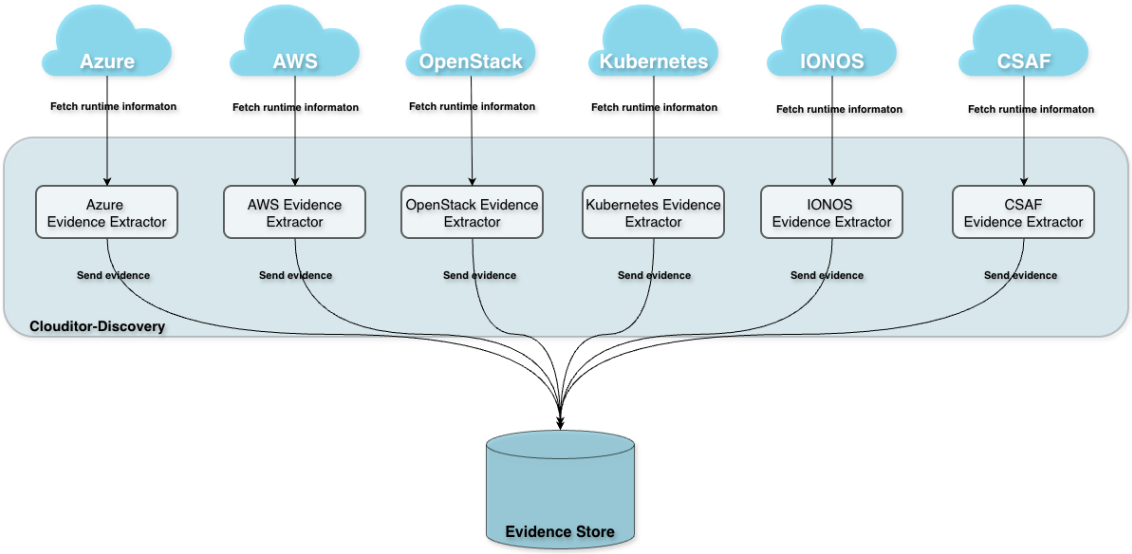

Figure 2. Overview of the Clouditor-Discovery component to extract runtime information from the Cloud Providers and submitting it as evidence to the Evidence Store in the EMERALD framework

For detailed interface specification, refer to the `./api/discovery` folder within the *Clouditor*[13] repository. The specification is written in accordance with the *Protocol Buffer Version 3 Language Specification*[14]. In addition, an auto-generated *.yaml* file that complies with the OpenAPI description for REST APIs is available in the `./openapi/discovery` folder. The implemented API functions are shown in Table 4.

Table 4. Overview of the important API functions for the Clouditor-Discovery

Function Name	Parameters	Return type	Description
`Start`	optional `resource_group` (string) optional `csaf_domain` (string)	successful (bool)	Triggers the start of the discovery process. Returns true if the component started without errors. The function includes two optional parameters: • With the parameter `resource_group` only the specified resource group is discovered. • With the parameter `csaf_domain` only the specified CSAF[15] domain is discovered.

3.2.1.1 Component description

The functionality of the *Clouditor-Discovery* component can be divided into 3 parts:

[11] https://go.dev/
[12] https://grpc.io/
[13] https://github.com/clouditor/clouditor/
[14] https://protobuf.dev/programming-guides/proto3/
[15] https://docs.oasis-open.org/csaf/csaf/v2.0/os/csaf-v2.0-os.html

- Fetching security-relevant configurations of cloud resources
- Creation of evidence based on the EMERALD evidence format (*CertGraph* Ontology)
- Forwarding the evidence to the *Clouditor-Evidence Store*

In *Clouditor-Discovery*, the discovery package is located at the top level. Its primary role is to communicate with the *Evidence Store component* in EMERALD. Initially, this service connects to the *Evidence Store*, then activates different discoverers (e.g., a discoverer for Azure) and continuously sends the gathered evidence through a gRPC channel to the *Evidence Store*.

For each CSP there is a separate sub-package (e.g., AWS, Azure, Kubernetes, OpenStack). Each package contains a CSP-specific file which initializes configurations and credentials that all underlying cloud services share. For the main resource types (e.g., compute, network, storage) there are corresponding Go files that fetch the desired runtime information from the service via API calls. According to the *CertGraph Ontology*, these properties are converted to the EMERALD evidence format which is independent from the used CSP. By using the tool *Owl2proto* the *Clouditor-Discovery* can directly use the automatically generated proto files without having to manage the creation and updating of the Ontology objects. The properties available for retrieval depend on the range of API calls offered by the respective CSP.

3.2.2 Technical specifications

The *Clouditor-Discovery* component is developed in Go (version 1.22) and consists of six packages for the supported providers. The following list outlines the supported providers along with their corresponding key modules:

- AWS: `github.com/aws/aws-sdk-go-v2`
- Azure: `github.com/Azure/azure-sdk-for-go`
- CSAF: `github.com/csaf-poc/csaf_distribution/v3`
- Kubernetes: `k8s.io/api`
- `OpenStack: github.com/gophercloud/gophercloud/v2`
- `IONOS: github.com/ionos-cloud/sdk-go/v6`

A full list of used libraries can be found in the *Clouditor* GitHub repository[16].

3.3 Delivery and usage

The following sections give a short overview of the delivery and usage of the prototype. Further technical details can be found in the *Clouditor* GitHub repository[17].

3.3.1 Package information

Table 5 shows the structure of the important folders of *Clouditor-Discovery* and a brief description of them.

[16] https://github.com/clouditor/clouditor/tree/main/service/discovery
[17] https://github.com/clouditor/clouditor

Table 5. Overview of the package structure of Clouditor-Discovery.

Folder	Description
`api/discovery`	This folder contains the interface specification file (.proto), the auto-generated *Protobuf,* gRPC files and the code needed for the communication. The .proto file defines the data structure and service methods using Protocol Buffers (Protobuf), a serialization format developed by Google for efficient data exchange.
`api/ontology`	This folder contains the auto-generated *.proto* file based on the *CertGraph* Ontology. The *.proto* file is generated by the *Owl2proto* tool, which converts Ontology files to Protobuf. For further information see D2.1 [12].
`cli/commands/service/discovery`	This folder contains the *Clouditor-Discovery* CLI-based source code files.
`cmd/discovery`	This folder contains the main files for the *Clouditor-Discovery* component.
`openapi/discovery`	This folder contains the auto-generated OpenAPI files.
`server/commands/discovery`	This folder contains the CLI command to start the *Clouditor-Discovery* server.
`rest/`	This folder contains the REST gateway implementation.
`service/discovery`	This folder contains the source code for the *Clouditor-Discovery* component separated for each CSP: AWS, Azure, Kubernetes, OpenStack, IONOS, and CSAF.

3.3.2 Installation

The installation instructions can be found in the *Clouditor* README[17]. If only the *Clouditor-Discovery* component should be started, the build and start command must be adapted to

```
go build -o ./clouditor-discovery cmd/discovery/discovery.go
```
and

```
./clouditor-discovery
```

Using the `--help` flag displays all available options as shown in Figure 3.

```
$ ./discovery  --help
This command starts a Clouditor Discovery service

Usage:
  discovery [flags]

Flags:
      --api-cors-allowed-headers stringArray   Specifies the headers allowed in CORS (default [Content-Type,Accept,Authorization])
      --api-cors-allowed-methods stringArray   Specifies the methods allowed in CORS (default [GET,POST,PUT,DELETE])
      --api-cors-allowed-origins stringArray   Specifies the origins allowed in CORS
      --api-default-password string            Specifies the default API password (default "clouditor")
      --api-default-user string                Specifies the default API username (default "clouditor")
      --api-grpc-port uint16                   Specifies the port used for the Clouditor gRPC API (default 9091)
      --api-http-port uint16                   Specifies the port used for the Clouditor HTTP API (default 8081)
      --api-jwks-url string                    Specifies the JWKS URL used to verify authentication tokens in the gRPC and
                                               HTTP API (default "http://localhost:8080/v1/auth/certs")
      --api-key-password string                Specifies the password used to protect the API private key (default "changeme")
      --api-key-path string                    Specifies the location of the API private key (default "~/.clouditor/api.key")
      --api-key-save-on-create                 Specifies whether the API key should be saved on creation. It will only created if
                                               the default location is used. (default true)
      --assessment-url string                  Specifies the Assessment URL (default "localhost:9093")
      --cloud-service-id string                Specifies the Cloud Service ID (default "00000000-0000-0000-0000-000000000000")
      --dashboard-callback-url string          The callback URL of the Clouditor Dashboard. If the embedded server is used, a public
                                               OAuth 2.0 client based on this URL will be added (default "http://localhost:8080/callback")
      --db-host string                         Provides address of database (default "localhost")
      --db-in-memory                           Uses an in-memory database which is not persisted at all
      --db-name string                         Provides name of database (default "postgres")
      --db-password string                     Provides password of database (default "postgres")
      --db-port uint16                         Provides port for database (default 5432)
      --db-ssl-mode string                     The SSL mode for the database (default "disable")
      --db-user-name string                    Provides user name of database (default "postgres")
      --discovery-auto-start                   Automatically start the discovery when engine starts
      --discovery-csaf-domain string           The domain to look for a CSAF provider, if the CSAF discovery is enabled
  -p, --discovery-provider strings             Providers to discover, separated by comma
      --discovery-resource-group string        Limit the scope of the discovery to a resource group (currently only used in the Azure discoverer
  -h, --help                                   help for discovery
      --log-level string                       The default log level (default "info")
      --service-oauth2-client-id string        Specifies the OAuth 2.0 client ID (default "clouditor")
      --service-oauth2-client-secret string    Specifies the OAuth 2.0 client secret (default "clouditor")
      --service-oauth2-token-endpoint string   Specifies the OAuth 2.0 token endpoint (default "http://localhost:8080/v1/auth/token")
```

Figure 3. Overview of the available options for Clouditor-Discovery

3.3.3 Instructions for use

Within the EMERALD project, the *EMERALD UI* is developed and used to access and manage the workflow within the framework. The *Clouditor-Discovery* is not directly accessible through the *EMERALD UI*; however, it can be attached to a Target of Evaluation in the EMERALD UI.

Figure 4 shows the *EMERALD UI* view for selecting an evidence extractor.

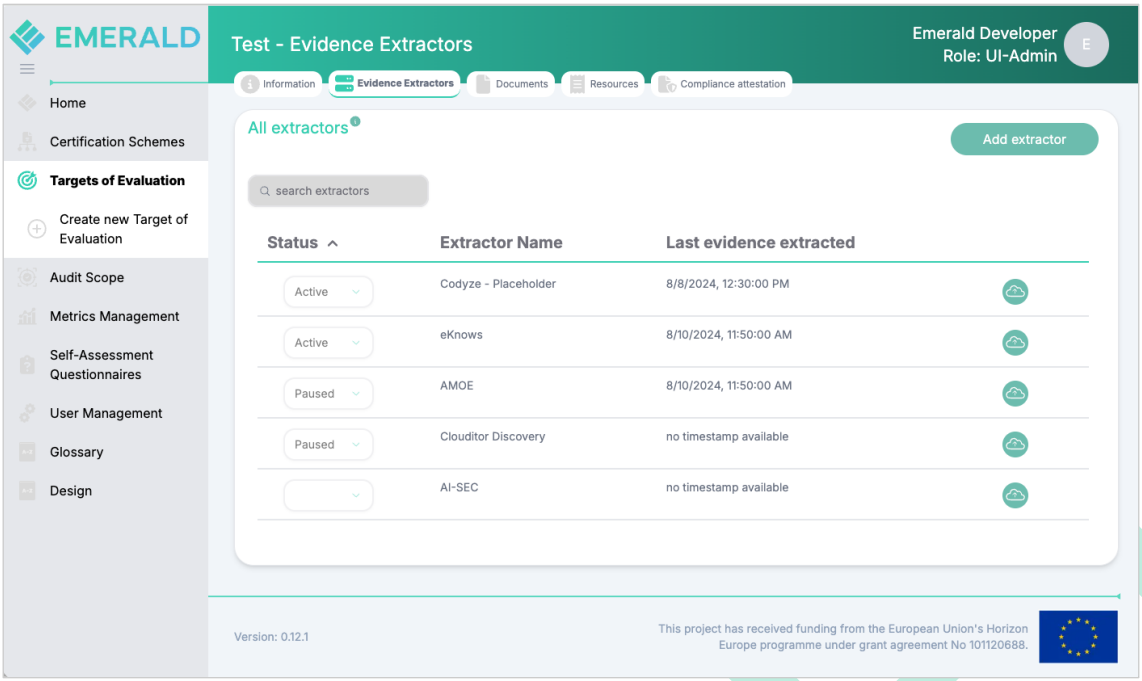

Figure 4. EMERALD UI view for setting the evidence extractor for a Target of Evaluation

Figure 5 presents the subsequent view containing configuration information for the available evidence extractors, in particular *Clouditor-Discovery*.

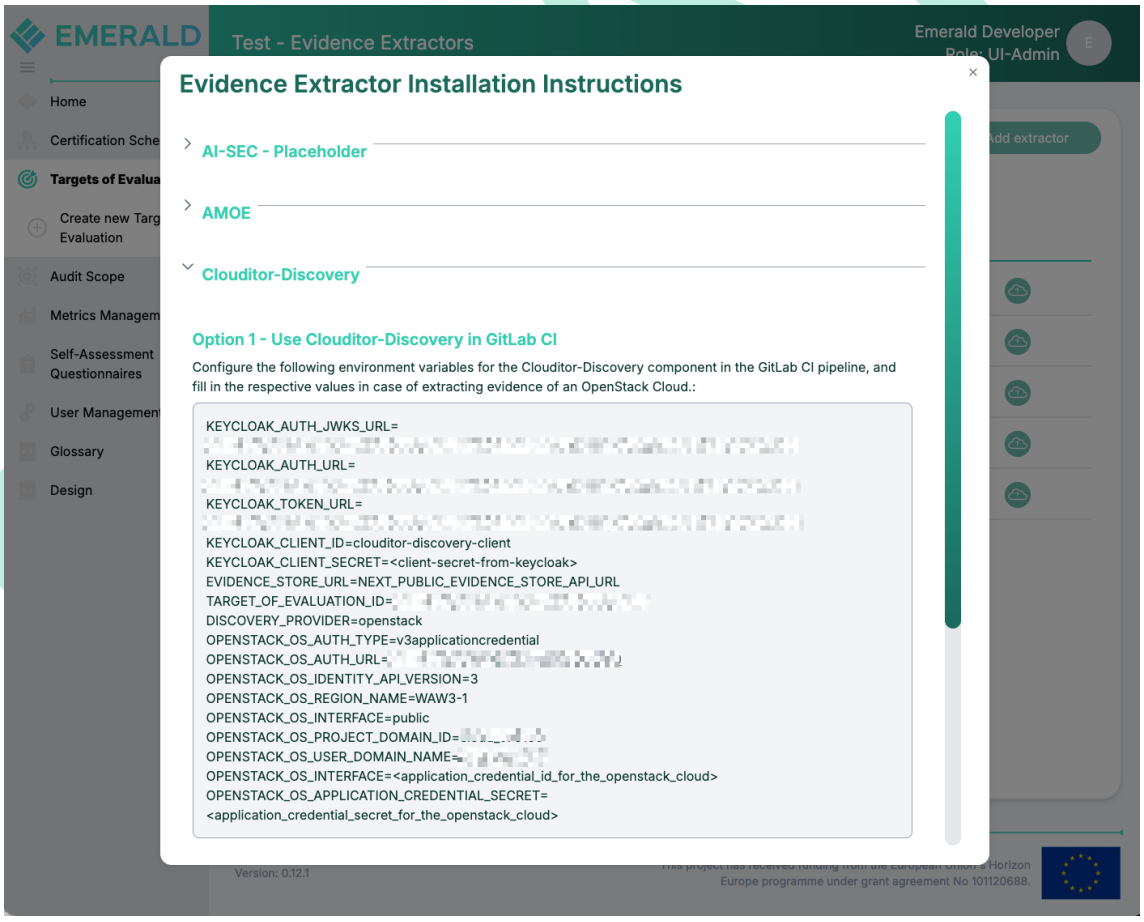

Figure 5. EMERALD UI view for the available evidence extractors and their installation instructions

Further information regarding the *EMERALD UI* can be found in D4.4 [13] and D4.6 [14].

Note that the installation of the evidence extractors cannot be performed via the *EMERALD UI*. The extractors can be configured in the GitLab CI pipeline (see Figure 5) or as described in the corresponding README[18]. The corresponding *Target of Evaluation* must be configured via the environment variables.

The available user manual of *Clouditor-Discovery* can be found in the *Clouditor* README[19] and the CLI[20].

3.3.4 Licensing information

The *Clouditor* is licensed under the open-source Apache-2.0 license including all sub-components such as *Clouditor-Discovery*.

3.3.5 Download

The *Clouditor* source code can be found in the *Clouditor* GitHub repository[21]. The adapted EMERALD component *Clouditor-Discovery* can be found in the public EMERALD GitLab repository[22].

3.4 Limitations and future work

Clouditor-Discovery currently gathers data from Microsoft Azure, AWS, OpenStack, IONOS, Kubernetes environments and CSAF, but its functionality is restricted by the access permissions granted to it in user management systems like *Azure Active Directory*. As a result, it can only access resources that are visible to the assigned user.

Changes to the cloud provider APIs may require updates to the component. For example, if key security features, such as at rest encryption, are modified, their integration into the evidence must be adjusted accordingly.

The evidence collection is restricted by the capabilities of the cloud provider APIs. If a specific encryption feature is not supported by an API, capturing evidence for that feature will not be possible.

Clouditor-Discovery incorporates ontological terms into the evidence; therefore, the constraints of the ontology must be considered:

- First, it is crucial that the ontology terms are accurately integrated into the evidence; otherwise, the assessment component may yield incorrect metrics.
- Second, the ontology requires ongoing updates, and any modifications must be reflected as well in the *Clouditor-Discovery* component.
- We have already developed a tool called *owl2proto* that converts the modelled ontology into an appropriate Protobuf schema, which can be directly used in various programming languages.

[18] https://git.code.tecnalia.dev/emerald/public/components/discovery
[19] https://github.com/clouditor/clouditor
[20] https://github.com/clouditor/clouditor#clouditor-cli
[21] https://github.com/clouditor/clouditor/tree/main/service/discovery
[22] https://git.code.tecnalia.dev/emerald/public/components/discovery

4 Codyze-Provenance

Codyze-Provenance is a new addition to the *Codyze* component in EMERALD (described in D2.3 [1]. It adds runtime evidence extraction capabilities to *Codyze* by creating a verifiable trail of evidence from source code to running cloud services and applications. This provenance creates a stronger link between static analysis based on source evidence extractors and runtime evidence extractors.

4.1 Functional description

Overall purpose. *Codyze-Provenance* utilizes the SLSA provenance framework[23] and the in-toto attestation framework[24]. The SLSA provenance framework specifies measures to harden the security of supply chains for software artefacts. In addition to specifying how each step in a software's supply chain can be security hardened, it also links the different stages to one another to support traceability and provenance. The in-toto attestation framework provides the technical specification on how to create these links by generating verifiable statements about what files went into a stage, what happened in a stage and what resulted from a stage. *Codyze-Provenance* uses provenance reports and attestations of software build processes to generate evidence. This information allows to track and verify what data went into a build process, what tools where executed and what artefacts were created. Thus, it becomes possible to trace the origin of a running cloud service and application back to its specific source code and corresponding build process.

Codyze-Provenance provides the necessary tool and supporting components to utilize SLSA and in-toto in a CI/CD pipeline. It submits provenance reports and attestations to the *Evidence Store* for further assessment. Finally, it also submits collected evidence to the *Trustworthiness System* as an alternative path to validate the authenticity of evidence.

Context and scope. *Codyze-Provenance* needs to be integrated into a CI/CD pipeline and executed on every run of the CI/CD pipeline producing a release for a cloud service or an application. It collects the provenance reports and attestations generated during a build. The collected information is transformed into evidence within the EMERALD framework and submitted to the *Evidence Store* and the *Trustworthiness System* [6].

Motivation. One challenge during runtime is to verify the origin of a running cloud service or application. Usually, it is possible to identify the container image or binary that is running. However, it becomes increasingly difficult to provide details on how the service or application was built, which built steps were executed, or what source code version was used in the build. A complete software supply chain would describe all resources and process producing the final deployable service or application. This provenance allows to verify requirements such as mandatory security testing of applications or security checks for dependencies. Moreover, a strong link between sources and build processes on the one hand, and the runtime on the other hand is created.

Requirements. Currently, *Codyze-Provenance* is a new addition in EMERALD. Its technical requirement with its implementation progress (degree of implementation - *partially / fully /not implemented*) is presented in Table 6.

[23] https://slsa.dev/
[24] https://in-toto.io/

Table 6. Requirement CODYZE.03 - Provenance reports from CI/CD

Field	Description
Requirement ID	CODYZE.03
Short title	Generate evidence from provenance reports.
Description	*Codyze* is executed as part of a CI/CD process. To better trace inputs like source code and outputs like build artefacts and link them to *Codyze* analysis reports, provenance reports are generated. These reports are collected and submitted as evidence.
Status	Validated
Priority	Could
Component	*Codyze*
Source	KPI
Type	Technical
Related KR	KR1_EXTRACT
Related KPI	KPI 1.1
Validation acceptance criteria	Validated if evidence arrives at the *Evidence Store* and the *Trustworthiness System*.
Progress	100% - Fully implemented
Milestone	MS6: Integrated audit suite V2 (M30)

Innovation. *Codyze-Provenance* provides a user-friendly approach to utilize SLSA and in-toto in software build processes for cloud services and applications. Collected evidence enhance the ability to assess compliance with respect to compliance schemes.

4.2 Technical description

The following subsections outline the technical details of *Codyze-Provenance*.

4.2.1 Prototype architecture

Codyze-Provenance consists of an application that collects provenance and attestation reports defined by the SLSA and in-toto framework. These reports are transformed into evidence for EMERALD and submitted to the *Evidence Store* and the *Trustworthiness System*. In addition, *Codyze-Provenance* provides templates and supporting files to create a CI/CD pipeline that supports the collection of provenance and attestation reports. Figure 6 depicts the high-level architecture of *Codyze-Provenance*.

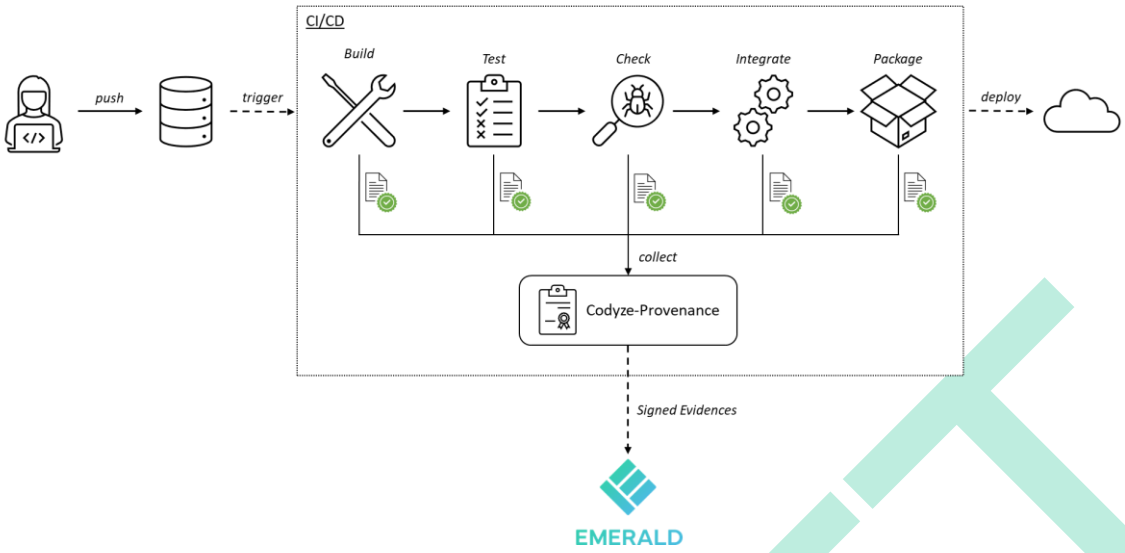

Figure 6. High-level architecture of Codyze-Provenance collecting provenance reports from a CI/CD pipeline and submitting it as evidence to the EMERALD framework (Evidence Store and Trustworthiness System)

4.2.1.1 Sub-components description

Codyze-Provenance has no sub-components.

4.2.2 Technical specifications

Codyze-Provenance is developed in the programming language Kotlin with a Java Virtual Machine as backend. Moreover, it provides templates and other supporting files to facilitate an integration into CI/CD pipelines. To this end, the focus is on GitLab, which is also used as the repository and CI/CD platform in EMERALD.

4.3 Delivery and usage

The following subsections detail the delivery and usage of *Codyze-Provenance*.

4.3.1 Package information

Codyze-Provenance is delivered as an application bundled in an archive. The structure of this package is summarized in Table 7.

Table 7. Overview of package structure for Codyze-Provenance

Folder	Description
`bin/`	Contains execution scripts for Windows and Linux/macOS
`docs/`	Contains detailed documentation
`etc/`	Contains sample configuration files and supporting files
`lib/`	Contains application and dependent libraries
`LICENSE`	License text (Apache License, Version 2.0)
`README.md`	Short documentation including short summary description, installation and usage instructions, and further information

Moreover, *Codyze-Provenance* is packaged as a container image for easier integration into CI/CD pipelines.

4.3.2 Installation

Installation instructions are provided as part of a README and documentation for *Codyze-Provenance* (cf. Table 7).

4.3.3 Instructions for use

Instructions on how to use *Codyze-Provenance* are provided as part of the released packages (cf. folder `'docs/'` and `'README.md'` in Table 7). This information is also available in the public GitLab repository[25] of EMERALD.

4.3.4 Licensing information

Codyze-Provenance is licensed as open source under Apache License, Version 2.0. In addition, it is ensured that third-party dependencies are compatible with the Apache License, Version 2.0.

4.3.5 Download

Codyze-Provenance is available from the public EMERALD GitLab repository[25] hosted by TECNALIA. The repository hosts the source code, the documentation, binary artefacts, and supporting materials.

4.4 Limitations and future work

Codyze-Providence is a new component added to the EMERALD framework. It collects and processes SLSA provenance reports and in-toto attestations into EMERALD evidence. It supports its users through its GitLab integration with the configuration of SLSA and in-toto. However, its main limitation is the complexity and diversity of CI/CD pipelines, which requires assistance from its users. They have to support the SLSA framework in their CI/CD pipelines and potentially adjust their build steps to generate reports and attestations.

[25] https://git.code.tecnalia.dev/emerald/public/components/codyze/codyze-provenance

5 Conclusions

The EMERALD project proposes a holistic approach to evidence collection, encompassing all levels of the cloud service from the infrastructure layer (e.g., virtual resources) to the business layer (e.g., policies and procedures), and the implementation layer (e.g., source code files).

This deliverable contains updates to deliverable D2.8 [7], an overview of the technological advances from the previous MEDINA project, and the technical report on the design, architecture, and current implementation status of the runtime evidence extractor components of Task 2.5 ("Extraction of evidence using runtime information"). The components adhere to the overall EMERALD framework and align with the technical requirements gathered within the scope of WP1. The deliverable outlines the relationship of the presented components with other components of the EMERALD framework and details the internal structure, subcomponents, and technical implementation information of each component.

The components introduced in this deliverable include the runtime evidence extraction components *Clouditor-Discovery* and *Codyze-Provenance*. The *Clouditor-Discovery* component was initially developed in MEDINA and has been further advanced to TRL 6, with a working prototype already integrated into the EMERALD framework. It has been enhanced to incorporate a variety of discovered resources and the use of the developed *Owl2proto* tool for the automatic generation of the necessary Ontology objects. *Codyze-Provenance* is a new addition to the *Codyze* component, which extracts evidence from CI/CD pipelines, provides attestations and provenances for software builds and artefacts, and aims at a TRL 5-6.

6 References

[1] EMERALD Consortium, “D2.3 Source Evidence Extractor - v2,” 2026.

[2] EMERALD Consortium, “D2.5 AMOE - v2,” 2026.

[3] EMERALD Consortium, “D2.7 ML model certification - v2,” 2026.

[4] EMERALD Consortium, “D2.10 Certification Graph–v1: Integration of the graph with semantically linked and combined evidence,” 2025.

[5] EMERALD Consortium, “D2.11 Certification Graph– v2: Integration of the graph with semantically linked and combined evidence,” 2026.

[6] EMERALD Consortium, “D3.4 Evidence assessment and Certification - Implementation - v2,” 2025.

[7] EMERALD Consortium, “D2.8 Runtime Evidence Extractor - v1,” 2024.

[8] C. Banse, A. Schneider and I. Kunz, “owl2proto: Enabling Semantic Processing in Modern Cloud Micro-Services,” *To appear in: 16th International Conference on Knowledge Engineering and Ontology Development,* 2024.

[9] EMERALD Consortium, “D1.4 EMERALD solution architecture - v2,” 2025.

[10] EMERALD Consortium, “D1.2 Data modelling and interaction mechanisms-v2,” 2025.

[11] MEDINA Consortium, “D3.6 Tools and techniques for collecting evidence of technical and organisational measures-v3 (https://medina-project.eu/public-deliverables/),” 2023.

[12] EMERALD Consortium, “D2.1 Graph Ontology for Evidence Storage: Description of a uniform schema for storing and linking heterogenous data,” 2024.

[13] EMERALD Consortium, “D4.4 User interaction and user experience concept - v2,” 2025.

[14] EMERALD Consortium, “D4.6 EMERALD UI - v2,” 2026.