



# DiSSCo Transition Project

## **Minimum Viable Product of Core Infrastructure**

### Work Package 3 - Deliverable 3.1

Wouter Addink (Naturalis),  
Sam Leeflang (Naturalis),  
Soulaine Theocharides (Naturalis)



Grant number 101130121

HORIZON-INFRA-2023-DEV-01-02 — Early phase implementation of ESFRI Projects  
which entered the ESFRI Roadmap in 2018

HORIZON Coordination and Support Actions - Horizon Europe

**27 May 2025**

**Version 1.1**

## Table of contents

<b>List of abbreviations</b>	<b>6</b>
<b>1. Introduction</b>	<b>7</b>
<b>2. Objectives</b>	<b>8</b>
<b>3. Development process</b>	<b>8</b>
<b>4. Machine Annotation Services support</b>	<b>9</b>
<b>5. Delivered product components</b>	<b>10</b>
5.1. Digital Specimen Repository	11
5.2. Persistent Identifier Infrastructure	11
5.3. Data processing and publishing	12
5.4. Authorisation and Authentication infrastructure (AAI)	13
5.5. Indexing and API	13
5.6. Annotation platform (DiSSCover)	13
5.7. Documentation	13
<b>6. Impact</b>	<b>14</b>
<b>7. Discussion</b>	<b>15</b>
<b>8. References</b>	<b>15</b>

**Project** DiSSCo Transition**Grant Agreement n°** 101130121**Start of the project** 01 December 2023**Duration** 18 months**Project coordinator** Stichting Naturalis Biodiversity Center**Deliverable title** Minimum Viable Product of Core  
Infrastructure**Deliverable n°** 3.1**Type of deliverable** Other**Work Package** 3**Lead Beneficiary** Stichting Naturalis Biodiversity Center**Dissemination level** PU - PublicAddink, W., Leeflang, S., & Theocharides,  
S. (2025). Minimum Viable Product of**Citation** Core Infrastructure - Deliverable 3.1  
(v1.1). DiSSCo Transition Project.<https://doi.org/10.5281/zenodo.15230656>**Due date of deliverable** 31 May 2025**Actual submission date** 27 May 2025**Quality review** Yes**Deliverable status**

Version	Status	Date	Responsible	Actions
1.0	Draft 1.0	29 April 2025	Wouter Addink	Sent for review
1.0	Draft 1.0	15 May 2025	Claus Weiland, Laurence Livermore	Reviewed
1.0	Final 1.1	26 May 2025	Wouter Addink	Finalised, with incorporation of feedback from reviewers
1.1	Final 1.1	26 May 2025	Wouter Addink	Sent for Coordinator verification
1.1	Final 1.1	26 May 2025	Vânia Ferreira	Verified

1.1	Final	27 May 2025	Wouter Addink	Submitted
-----	-------	-------------	---------------	-----------

**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 or the European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

## Executive Summary

This document outlines the Minimum Viable Product (MVP) of the DiSSCo core infrastructure, developed by Naturalis to support digital specimens and enable services like machine annotation. The MVP advances the infrastructure from Technology Readiness Level (TRL) 4-5 to TRL 6, demonstrating its use with real data. It was delivered ahead of schedule to support a hackathon focused on annotation services.

The MVP includes production-ready versions of key infrastructure components (collectively referred to as DSArch) that provide:

- Persistent identifiers (DOIs and Handles)
- Metadata harmonisation via the openDS data model
- Annotation support, including machine-actionable FAIR Digital Objects
- Provenance tracking and versioning
- REST APIs for data interaction

A sandbox environment for testing, and documentation were also made available.

The MVP supports human and machine annotations using a standardized format, with machine annotation services operating via middleware in the infrastructure, pending formal Service Level Agreements (SLAs) with DiSSCo ERIC for production deployment.

The MVP emphasises community involvement, with features defined through consultations and stakeholder workshops. It promotes data reuse, interoperability, and the use of DOIs in scholarly referencing of specimens. Currently, 9 Machine Annotation Services (MASs) using the infrastructure are in sandbox testing, including tools for label digitisation and linking data.

Impact areas include improved specimen referencing in research, industrialisation of digitisation, and better use of taxonomic expertise. Feedback from hackathons has informed the roadmap for further development, especially in areas like image handling for MASs.

## Keywords

DiSSCo, FAIR Digital Objects, Annotations, Machine Actionable, Data Infrastructure, Digital Specimens, Digital Extended Specimen

## List of abbreviations

AAI - Authorisation and Authentication Infrastructure

AI - Artificial Intelligence

CA - Consortium Agreement

CI/CD tools - Continuous Integration/ Continuous Development tools

D - Deliverable

DiSSCo - Distributed System of Scientific Collections

DOI - Digital Object Identifier, a globally unique, resolvable and persistent identifier

DTP - DiSSCo Transition Project

DwC - Darwin Core data standard

ERIC - European Research Infrastructure Consortium, a legal entity for European RIs

FAIR - Findable, Accessible, Interoperable and Reproducible, or Fully AI Ready

FDO - FAIR Digital Object, a structured and technology-agnostic digital package that combines data, metadata, and persistent identifiers to be machine-actionable

GA - Grant Agreement

GUI - Graphical User Interface; enables a person to interact with digital services

iGA - Interim General Assembly, a body to support the preparation of the RI consisting of DiSSCo partner institutions

IPDES - International Partners for the Digital Extended Specimen

MAS - Machine Annotation Service, a service that provides automated annotations

MIDS - Minimum Information about a Digital Specimen.

MS - Milestone

PID - Persistent Identifier

RA - DOI Registration Authority

RI - Research Infrastructure

RO - Research Output

SLA - Service Level Agreement

TRL - Technology Readiness Level, an indicator for maturity of technology

WP - Work Package

## 1. Introduction

This document describes the Minimum Viable Product (MVP) of core infrastructure delivered by Naturalis as delivered in the project. The MVP is a production version of core infrastructure components that have been prototyped in earlier projects.

The MVP was delivered ahead of the planning, so that a hackathon, organised by Senckenberg, could be organised during the project to develop additional machine annotation services in a sandbox environment (a copy of the core infrastructure with test data). The functionality to include in the MVP was based on consultations with the DiSSCo Technical Advisory Board (independent experts<sup>1</sup>) and with stakeholders: the National Nodes, the Consortium of European Taxonomic Facilities–Information Science<sup>2</sup> and Technology Committee (CETAF ISTC<sup>3</sup>) working group and project partners in WP3 and WP4.

Core infrastructure components that have been put in production are described in chapter 4. These are the components as defined in the project description, plus a portal called DiSSCover that enables users to make annotations (Koureas 2024). In DiSSCo documentation, the core infrastructure is also named DSArch: Digital Specimen Architecture. The core infrastructure plus source systems (e.g. collection management systems) provided by the DiSSCo partners together form the distributed data infrastructure in DiSSCo. The core infrastructure provides APIs to enable DiSSCo services and other services to interact with the digital specimen data, see Figure 1. The infrastructure acts as a platform that harmonises data from DiSSCo facilities, assigns persistent identifiers, records provenance, supports linking and extending the data and supports annotations. Enriched data and enhanced data can be consumed by, e.g. other RIs or by the DiSSCo facility source systems, and DiSSCo services can exchange data with the platform.

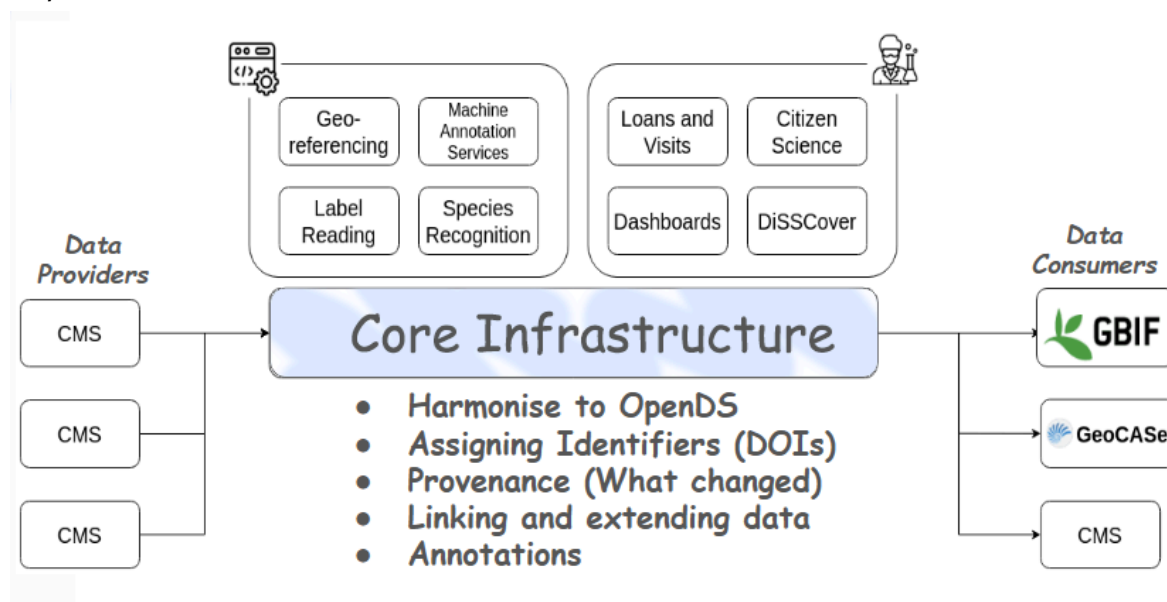


Figure 1: Core infrastructure as a platform for DiSSCo services.

<sup>1</sup> <https://www.dissco.eu/dissco/governance>

<sup>2</sup> <https://cetaf.org>

<sup>3</sup> <https://istc.cetaf.org>

This document includes a short description of the developed components but does not go into detail about the underlying concepts of the core infrastructure, such as the FAIR Digital Object concept, the Digital (Extended) Specimen concept or the concept of machine actionability. More information about underlying concepts is available in earlier architectural documents—particularly relevant are DPP D6.2 Implementation and construction plan of the core architecture<sup>4</sup>, Architectural overview documents<sup>5</sup> and the RFC documents<sup>6</sup>.

## 2. Objectives

The aims of the MVP were described in Milestone 14:

1. Advance from TRL 4-5 (proof of concept tested with Handle infrastructure) to TRL 6 (scale as used in EU/Horizon projects<sup>7</sup>): Technology demonstrated in a relevant environment with a representative set of real data.
2. Provide a minimum set of features that show clear added value of the DiSSCo core infrastructure to attract early-adopter users of the RI.

For the second aim, the focus was on enabling annotation services. The infrastructure has been designed to support Digital Specimens as mutable, versioned objects that can act as a surrogate for the physical specimen by providing access to all known information about the specimen that is digitally available. These objects are designed to be community curated, and full provenance is recorded. The objects are also designed to be "Fully AI-Ready" by implementing them as machine-actionable FAIR Digital Objects<sup>8</sup>. Objects have machine-readable descriptions of context (metadata) and data structure, are identified by persistent identifiers and can be annotated by machines that use AI. The MVP enables annotations by both humans and machine annotation services, and already includes annotations that showcase this.

Besides support for annotations and provenance data, the MVP also provides added value by harmonizing specimen data to the openDS data model<sup>9</sup> (supporting community standards such as [Darwin Core](#)), by providing persistent identifiers for all digital objects in the infrastructure and by providing an indicator for completeness of the specimen metadata ([MIDS](#) level, where the highest achievable level is MIDS 3).

## 3. Development process

Core infrastructure components were initially developed up to TRL 4-5 through funding received from DiSSCo partners (iGA members) and Naturalis in-kind contribution and through EU-funded projects: [SYNTHESYS+](#) (AAI components, Smith et al, 2019), [DiSSCo Transition](#) (openDS datamodel including the PID scheme, Hardisty 2021) and [BiCIKL](#)

<sup>4</sup> <https://doi.org/10.5281/zenodo.6832200>

<sup>5</sup> <https://drive.google.com/drive/folders/1YLeGKXf5WqlybbeFyGQoGRVWV9dMPZnW>

<sup>6</sup> <https://github.com/DiSSCo/openDS/issues?q=is%3Aissue+is%3Aclosed+rfc>

<sup>7</sup> Technology readiness levels (TRL); Extract from Part 19 - Commission Decision C(2014)4995" (PDF). ec.europa.eu. 2014.

<sup>8</sup> <https://fairdo.org/specifications/>

<sup>9</sup> <https://terms.dissco.tech> <https://cordis.europa.eu/project/id/101007492>



(Persistent Identifier Infrastructure and Digital Specimen Repository components). These were further developed into a production version that includes a first set of specimen data from DiSSCo facilities, provides DOIs through membership with DataCite, and includes documentation for users and DiSSCo service providers that want to build services on top of this infrastructure.

The development was carried out by a small DiSSCo development team at Naturalis: Sam Leeflang (lead developer), Souleine Theocharides (backend developer) and Tom Dijkema (frontend developer), with support from Sharif Islam (FAIR and Open Science Lead) and coordinated by Wouter Addink (DiSSCo deputy director).

The functionality to include in the MVP was based on consultations with the DiSSCo Technical Advisory Board (independent experts) and with stakeholders: the National Nodes, the CETAF ISTC working group and project partners in WP3 and WP4. This functionality is described in detail in the milestone document (Milestone 14). The software and documentation were developed in sprints (following agile development principles and using CI/CD tools), and progress was presented for discussion with stakeholders in bi-monthly online sessions. These sessions were recorded and can be accessed here: <https://dissco.tech/> and through the [DiSSCo RI YouTube channel](#). All software is published as open source through the [DiSSCo GitHub repository](#).

The openDS data models used in the MVP have been documented and undergone a public review process in which international stakeholders involved in development of biodiversity information standards, as well as stakeholders in IPDES and DiSSCo, have been asked to comment through a standardised review process (see <https://github.com/DiSSCo/openDS/issues> for received feedback). The reviewed version used in the MVP is 0.4 and can be found at [schema.dissco.tech](https://schema.dissco.tech).

For the MVP, three separate environments are provided: a production environment, an acceptance environment and a development environment. The GUI for these environments can be accessed through, respectively, <https://disscover.dissco.eu>, <https://sandbox.dissco.tech> and <https://dev.dissco.tech>. A full overview of all internet addresses in use for the MVP, including the APIs, can be found in the developer documentation (see [5.7](#)).

The sandbox contains test data with test Handles as identifiers instead of DOIs and is used for acceptance testing, as well as for testing machine annotation services. The test data and identifiers for tests will be deleted from time to time.

## 4. Machine Annotation Services support

The core infrastructure allows Machine Annotation Services (MASs) to operate through a piece of middleware that runs in the infrastructure. This is done so that such services can be scaled based on the requested load. The middleware is provided by the MAS service provider, a template for this is available (see [5.7](#) documentation). The middleware (also known as wrapper) calls a value service, an external service that operates as a black box

and provides the automated annotation. For this, the middleware supplies information about the object to annotate to the value service and supplies the generated annotation(s) back to the infrastructure. Annotations are created in a standard, W3C Web Annotation Data Model compliant format and are stored also as FDOs, so annotations can themselves also be annotated, e.g. by an agent validating the auto-generated annotation. MAS middleware software is peer-reviewed by developers from the core infrastructure host and undergoes automated code quality checks to ensure secure operation.

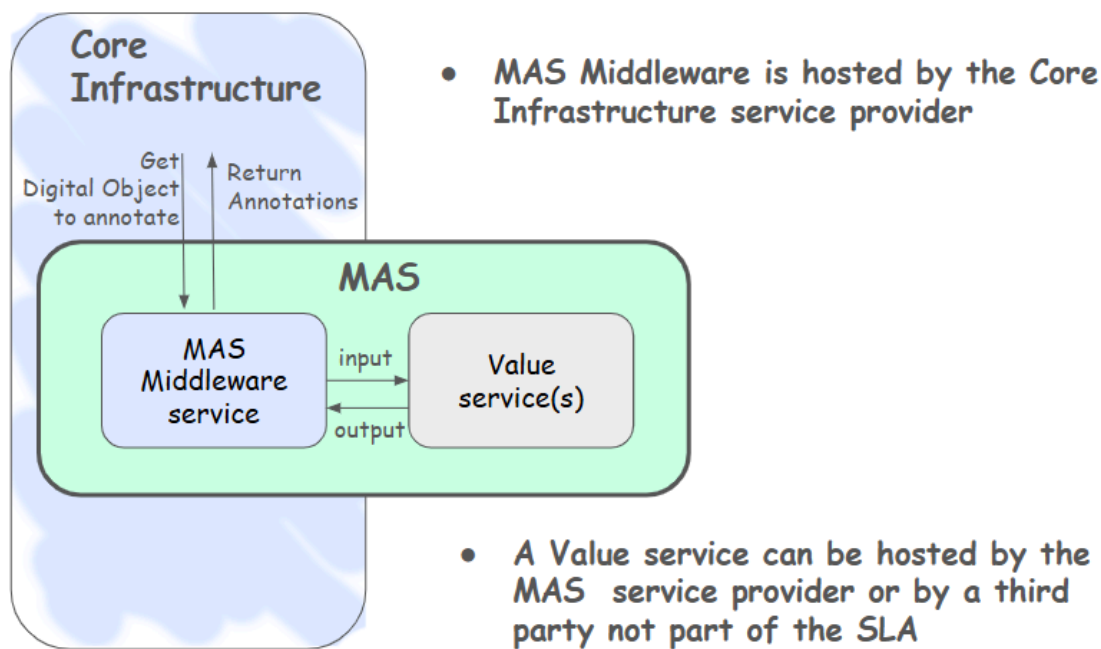


Figure 2: Core Infrastructure support for MASs

Machine annotation services are currently only allowed to run in a sandbox environment, as these are provided by third parties (service provider nodes, SPs) and a service level agreement (SLA) with DiSSCo ERIC is required to run these in the production environment. A SLA is a contractual agreement that can only be made when the ERIC is in place and is needed to ensure trustworthy, reliable services. In preparation for the ERIC, the DiSSCo Transition Project already developed a corpus of draft policies to underpin this, and also a General Framework Agreement for Service Providers was developed with stakeholders to form the basis for SLAs. Specifically for MASs, a draft MAS policy<sup>10</sup> and a SLA template<sup>11</sup> was created to provide potential MAS service providers with information about the requirements to operate these in DiSSCo. A SLA for MASs has specific requirements as it is an agreement between three parties: the Service Provider, Service Host (the provider of the core infrastructure, Naturalis) and the Service Consumer (DiSSCo ERIC). Also, a MAS often makes use of third-party services such as LLMs that are not under the control of the MAS service provider.

## 5. Delivered product components

<sup>10</sup> [Draft MAS policy](#)

<sup>11</sup> [SLA template for MAS](#)

## 5.1. Digital Specimen Repository

The repository stores digital objects in the core infrastructure. Digital object types provided in the MVP are digital specimens, digital media, annotations, create-update-tombstone event records (for provenance), machine annotation service descriptions, data mappings and source system descriptions. Both the latest version, as well as deltas, are stored to support versioned objects and indexing. The component is described in detail in the BiCIKL deliverable: Digital Object Interface Protocol [DOIP] enabled Digital Object repository installation to store and provide digital specimen information<sup>12</sup>. For the MVP, production and sandbox versions were created, and first datasets were ingested from Darwin Core and ABCD endpoints provided by DiSSCo facilities. The sandbox environment was extended with an extra option to include datasets with selected specimens of interest for MAS development as selected by MAS developers, to support their needs for testing MASs.

## 5.2. Persistent Identifier Infrastructure

The persistent identifier infrastructure is described in detail in BiCIKL deliverable 7.1 (Addink et. al, 2023). For the MVP, a DOI prefix (10.2020) was provided by DataCite (10.3535). Because DataCite is not capable (yet) of creating DOIs with metadata in the PID record as needed for FAIR Digital Object support, and because of current performance constraints in the DataCite infrastructure, the MVP uses its own PID infrastructure for both internal PIDs (for which Handles are generated), as for the DOIs.

After creation of a DOI, the infrastructure sends the DOI metadata asynchronously to DataCite so that it also becomes available in the DataCite metadata repository. The metadata is used in services like the DOI citation service. DataCite will evaluate what optimisation of their infrastructure is needed when it has received metadata from about 5 million DOIs created by DiSSCo. In case the scalability and performance of the DataCite infrastructure becomes a major bottleneck for DiSSCo, there are a few alternatives: sending metadata to DataCite without using the DataCite API, choosing another DOI registration authority or becoming a new registration authority under the DOI foundation. The global Handle infrastructure that is powering DOIs itself has no scalability issues in serving the billions of DOIs that DiSSCo plans to create.

An example DOI as created by DiSSCo MVP:

<https://doi.org/10.3535/VH4-78T-FMW>

To see the metadata in the PID record:

<https://doi.org/10.3535/VH4-78T-FMW?noredirect>

Note that the DOI redirects both to the human-readable page in DiSSCover (default) as well as to a machine-readable version of the digital object and to the online specimen catalogue record from the DiSSCo facility. A user can, for example, be redirected to the catalogue record view by suffixing it with the locatt attribute:

<https://doi.org/10.3535/VH4-78T-FMW?locatt=view:catalog>

<sup>12</sup> <https://cordis.europa.eu/project/id/101007492/results>

The DOI directs to the latest version of the digital object by default, but can also be used to go to a specific version by appending the version number to the DOI:

<https://doi.org/10.3535/VH4-78T-FMW?urlappend=/1>

A Digital Specimen DOI should be cited as follows:

Tallinn University of Technology (2025). sedimentary rocks. Distributed System of Scientific Collections. [Dataset]. <https://doi.org/10.3535/VH4-78T-FMW>

In HTML pages, the DOIs created by the MVP can provide contextual information from the PID record through a tooltip, as shown in Figure 3. This is already supported in the [Arpha journal system](#) (Pensoft).

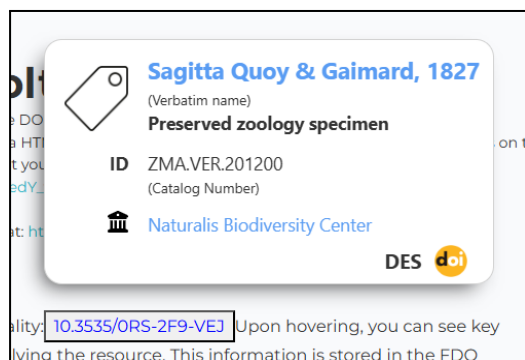


Figure 3: DOI tooltip

### 5.3. Data processing and publishing

The MVP is able to automatically process data from source systems (endpoints provided by DiSSCo facilities) that publish specimen data in Darwin Core archives or as ABCD(EFG) data through BioCAsE endpoints. The latter is used for biological specimens and earth science specimens, while Darwin Core is a standard for occurrence data, which at present does not support earth science specimens data. Mapping of a dataset to openDS can be customised, and these mappings, as well as descriptions of the source systems (endpoints), are stored in the digital object repository. An interface is available to support this: <https://orchestration.dissco.tech/>

In the sandbox environment, we support including datasets with selected specimens of interest for MAS development, as selected by MAS developers, to support their needs for testing MASs. MASs that are available in the sandbox are listed in the orchestration GUI for the acceptance environment:

<https://acc.orchestration.dissco.tech/>

To support the inclusion of Digital Specimen DOIs in the source systems, a generic download service was implemented. The first download option is an export that provides a download file with physical specimen identifiers together with the Digital Specimen DOIs for a selected source system. This file can then be used by the source system provider to import the DOIs. A new DarwinCore term (digitalSpecimenID) to support inclusion of

digital specimen DOIs in DarwinCore data sets is, at the time of writing this deliverable, under public review: <https://github.com/tdwg/dwc/issues/530>.

#### **5.4. Authorisation and Authentication infrastructure (AAI)**

AAI infrastructure for DiSSCo was developed by GRNet as a pilot in the Synthesys+ project (Addink 2020 & Leeftang 2023b). It allows users and machines to get authenticated. Users can authenticate through ORCID, their institutional account, or a Google account, and ORCID is used to identify people's contributions to the data.

The developed prototype is still used in the MVP, since GRNet has made it clear that it cannot provide support on a production version and that it is not interested in creating a future SLA with DiSSCo ERIC as an AAI service provider. Naturalis therefore created a new AAI service based on a KeyCloak instance provided by Naturalis. Users will not need to be migrated as the MVP does not yet have many users. This new service will soon replace the GRNet service and will provide similar capabilities but with an improved GUI. The DiSSCo CSO is discussing with NHM London if they are interested in becoming the SP for this service in the future.

Apart from the core infrastructure service provider, two other potential SPs are planning to use the AAI infrastructure in the near future to authenticate users: for the European Loans and Visits System and for the Specialisation Plan services.

#### **5.5. Indexing and API**

Data in the infrastructure is indexed through an Elasticsearch index and provided through a REST API using the JSON:API specification (Leeftang 2023a). An (experimental) DOIP API is also available, as described in BiCIKL's Deliverable 7.4 (Theocharides et al, 2025). APIs for production are available through <https://api.dissco.eu>, in the format [endpoint]/[version], for example <https://api.dissco.eu/digital-specimen/v1>. The API is used by, e.g. DiSSCover and ELViS to search and retrieve data. For the MVP, an extra API was created to provide digital specimen data to ELViS.

#### **5.6. Annotation platform (DiSSCover)**

While a GUI is not part of the core infrastructure, DiSSCover was included in the MVP as GUI to provide an easy way for users to use the annotation capabilities of the core infrastructure. DiSSCover is one of the core end-user services that DiSSCo aims to provide, which makes use of the core infrastructure.

#### **5.7. Documentation**

[Developer Documentation](#) for developers is maintained in DiSSCo's GitHub, which publicly describes development methods, release management and architecture. It also contains an overview of all domain names in use. [MAS developer documentation](#) has been provided through a dedicated documentation site and was used in the hackathon. The API

is documented through [Swagger documentation](#). There is also detailed developer [documentation for DiSSCover](#).

## 6. Impact

With the MVP of the Core Infrastructure in place, a workshop was organised together with Work Package 2 in which stakeholders could give input on the roadmap and prioritisation for further development of the Core Infrastructure. The next step is to support early adopters in using the infrastructure. User feedback from these stakeholders can then be used to improve the infrastructure, and early adopters will promote the Core Infrastructure further in the community when their needs can be satisfied. To create impact, the first focus will be on Annotation services, in particular MAS, and on the use of the Digital Specimen DOIs. The latter will require sufficient mass in digital specimens as well as support in community standards like Darwin Core. An additional term for this has been proposed and is in [public review](#), see also D4.1. Also, a download service has been developed in DiSSCover to support the download of DOIs for inclusion in DwC datasets and local collection management systems.

DiSSCo is working with DataCite to speed up the creation of DOIs, and the implementation of metadata in the DOI Handle record is fuelling discussions in the DOI Foundation, which may impact future handling of DOI metadata. Currently, metadata is usually stored in metadata repositories at the registration authorities, which are siloed and lack a common standard for DOI metadata. This limits usage and makes metadata migration between RAs difficult, which is a potential issue for the long-term availability of DOIs. Storing metadata in the PID record with machine-readable definitions, as implemented in DiSSCo, has the potential to solve such issues.

The DOIs have an enormous potential in improving the way specimens are referenced in scholarly publications. Using persistent and resolvable identifiers instead of catalogue numbers that are not stable and often not globally unique is a major improvement that makes it much easier for taxonomic experts and other researchers to repeat research and find the referenced material. And having these identifiers linking to a digital specimen that provides access to all knowledge digitally available about the specimen, instead of just a catalogue record, also contributes to that. While catalogue records today often also get globally unique and resolvable identifiers, these are often not fully persistent because they include semantics, and they lack standardised PID metadata.

The support for annotations, in particular for automated annotations on the European dataset, will impact the quality and completeness of the specimen data, allows for optimized use of scarce taxonomic expertise that is often not available within an institution holding the relevant collections, reduces the time needed for 'data wrangling' and has the potential to speed up digitisation, since all metadata extraction from specimen images can be automated. In-house developed AI solutions to improve or digitise data in institutions can now also benefit other institutions in the infrastructure.



Through a MAS hackathon organised in the DTP project and through a capacity building project for using AI for label digitisation (the [CESP](#) funded project [AI4Labels](#)), the MVP to date already has 9 MASs available in the sandbox environment for testing, including MASs to create links to related data in other infrastructures, MASs for research purposes like extracting information about plant organs from herbarium specimens, and a MAS to support specimen label digitisation. An overview of available MASs can be found in the sandbox orchestration interface: <https://acc.orchestration.dissco.tech/>. Operation in the production environment is not yet possible as this will need the establishment of service level agreements with the ERIC, for which the ERIC needs to be in place.

## 7. Discussion

MVP functionality was discussed with stakeholders during the bi-monthly demonstration sessions, and priorities for further development were discussed in a [dedicated workshop](#) (milestone 10). The stakeholders indicated that working with early adopters to improve functionality and foster community uptake is one of the highest priorities. This was approached through hackathons with potential MAS providers, discussions with potential service providers, and a pilot is planned to investigate how a digitisation project for South-East Asian butterflies can benefit from the MVP annotation services. Working together with end-users will provide valuable feedback for the improvement of the DiSSCover annotation platform.

The hackathons provided feedback for further improvements to support MASs. It was observed that MASs often work on the images and also generate new images, such as heatmaps. This needs to be better supported by providing storage for these new images in the infrastructure and by optimising access to the images for MASs. Instead of downloading and resizing each image by every MAS, it would be more efficient to provide central storage for downsized images and ways to process these in batches by a MAS, although this comes with a cost for the infrastructure. Findings were added to the backlog, and the roadmap for further development was adjusted.

## 8. References

Addink, W., Islam S., Koumantaros K. , Laskaris N. (2020). MS48 AAI Infrastructure Design. SYNTHESYS PLUS <https://doi.org/10.17605/OSF.IO/HKEBA>

Addink W, Islam S, Dillen M, Güntsch A, Theocharides S (2023) Deliverable D7.1 Architecture Design for a pan-European PID system for Digital Specimens. ARPHA Preprints. <https://doi.org/10.3897/arphapreprints.e107168>

Hardisty AR, Addink W, Glöckler F, Güntsch A, Islam S, Weiland C (2021) A choice of persistent identifier schemes for the Distributed System of Scientific Collections (DiSSCo). Research Ideas and Outcomes 7: e67379. <https://doi.org/10.3897/rio.7.e67379>

Koureas D, Livermore L, Addink W, Alonso E, Alonso JM, Casino A, Dusoulier F, Ferreira V, Grieb J, Groom Q, Islam S, Köljal U, Lymer G, Marhold K, Paleco C, Pijls S, Scory S, Scott B,

Weiland C, Worley K (2024) DiSSCo Transition Abridged Grant Proposal. Research Ideas and Outcomes 10: e118241. <https://doi.org/10.3897/rio.10.e118241>

Leeftang, S., Weiland, C., Grieb, J., Dillen, M., Islam, S., Fichtmueller, D., Addink, W., & Haston, E. (2022). DiSSCo Prepare Deliverable D6.2 Implementation and construction plan of the DiSSCo core architecture. DiSSCo Prepare. <https://doi.org/10.34960/50B9-KJ05>

Leeftang, S., Weiland, C., Islam, S., Dillen, M., & Theocharides, S. (2023a). DiSSCo Prepare Deliverable D6.3 - A generalised set of API specifications for interaction with the DiSSCo core architecture. DiSSCo Prepare. <https://doi.org/10.34960/HTZV-NW73>

Leeftang, S., Liampotis, N., Georgilakis, C. & Islam, S. (2023b). D6.2 Piloting Access Through an AAI infrastructure. SYNTHESYS PLUS.

Smith VS, Gorman K, Addink W, Arvanitidis C, Casino A, Dixey K, Dröge G, Groom Q, Haston EM, Hobern D, Knapp S, Koureas D, Livermore L, Seberg O (2019) SYNTHESYS+ Abridged Grant Proposal. Research Ideas and Outcomes 5: e46404. <https://doi.org/10.3897/rio.5.e46404>

Theocharides S, Leeftang S, Addink W, Islam S (2025) Digital Object Interface Protocol (DOIP) enabled Digital Object repository installation to store and provide digital specimen information. ARPHA Preprints. <https://doi.org/10.3897/arphapreprints.e157339>