



High-Performance Computing Skills Platform
and European Collaboration for Training

D3.1 – Pilot Deployment of the EuroHPC Training Platform

WP3: Deployment of EuroHPC Training Platform

Grant Agreement Number: 101136267



Funded by
the European Union



EuroHPC
Joint Undertaking

This project has received funding from the European High-Performance Computing Joint Undertaking under grant agreement No. 101136267.

Document Information

Deliverable Number	D3.1
Deliverable Name	Pilot Deployment of the EuroHPC Training Platform
Due Date	31/12/2024 (M12)
Deliverable Lead	KIFU
Authors	Zoltan Kiss, KIFU; Simon Wong, ICHEC
Responsible Author	Zoltan Kiss, KIFU
Keywords	HPC Training Portal, portal deployment, event management platform, HPC training events, HPC training materials
WP	WP3
Nature	Report
Dissemination Level	Public
Final Version Date	27/12/2024
Reviewed by	Kevin Lüdemann, GWDG
PMT Approval	27/12/2024

Funded by the European Union. Views and opinions expressed are however those of the author only and do not necessarily reflect those of the European Union or the European High-Performance Computing Joint Undertaking. Neither the European Union nor the granting authority can be held responsible for them.

Document History

Partner	Date	Comments	Version
ICHEC	15/11/2024	Skeleton version	0.1
KIFU	04/12/2024	Draft with technical contributions	0.5
ICHEC	10/12/2024	First draft for internal review	0.8
KIFU	17/12/2024	Minor revisions from feedback	1.0
ICHEC	27/12/2024	Final version with PMT approval	1.1

Executive Summary

The HPC SPECTRA project aims to develop a comprehensive EuroHPC Training Platform which forms a key part of the EuroHPC JU's training strategy. This platform is designed to foster high-performance computing (HPC) skills across Europe, providing a centralised hub for accessing and promoting HPC training resources. The project supports European training providers by facilitating the promotion of events and materials, while offering users a centralised access point for finding and discovering relevant training opportunities.

The training platform will align with existing European efforts by integration with the C2ISS (CASTIEL 2 Information Sharing System), which will become a central hub in dissemination of European HPC activities. Having determined the main architectural requirements, concepts and elements of the training platform earlier in the project, this deliverable describes the deployment phase of the platform. This includes integration of key components like Indico for event management and Zenodo for archiving training materials. Other elements of the platform include structured metadata and taxonomies for better searchability, seamless user experiences via Single Sign-On (SSO), and automation for data synchronisation through APIs.

The deployment phase focused on developing robust content structures, taxonomies, and integration tools. Tasks included configuring metadata in Drupal and Indico, integrating event management via Indico, automating data exchanges between systems, and establishing a long-term materials repository with Zenodo. Throughout this process, workflows and user journeys were conceived to cater for the diverse needs of training providers, attendees, and content managers. Previews of the training platform as deployed have also been shown.

As the platform is prepared to be launched in January 2025, an initial rollout phase will involve limited on-boarding of a few training providers, with further testing and feedback mechanisms in place to ensure a smooth user experience. The full release, anticipated by mid-2025, will aim to onboard all users and training providers, marking a significant milestone in advancing HPC training across Europe.

List of Abbreviations

API	Application Programming Interface
BSC	Barcelona Supercomputing Center
CASTIEL 2	Project to coordinate the activities of the NCCs and CoEs at the European level
CERN	European Council for Nuclear Research
CoE	Centre of Excellence
CMS	Content Management System
DOI	Digital Object Identifier
EuroCC	National Competence Centres in the framework of EuroHPC
GWGD	Gesellschaft fuer wissenschaftliche Datenverarbeitung mbH Goettingen
HPC	High-Performance Computing
HTTP	Hypertext Transfer Protocol
ICHEC	Irish Centre for High-End Computing
IdP	Identity Provider
JSON	JavaScript Object Notation
JU	(EuroHPC) Joint Undertaking
KIFU	Kormányzati Informatikai Fejlesztési Ügynökség (Governmental Information-Technology Development Agency, Hungary)
NCC	National Competence Centre
PMT	Project Management Team
PRACE	Partnership for Advanced Computing in Europe
SAML	Security Assertion Markup Language
SSO	Single Sign-On
URL	Uniform Resource Locator
WP	Work Package

List of Figures

Figure 1. High-level overview of the C2ISS portal structure. Development of the Training section (highlighted in orange) is the responsibility of the HPC SPECTRA project.....	9
Figure 2. Interplay between C2ISS portal and other training components/services.	10
Figure 3. Miro canvas showing conceptual workings and sketches of different sets of metadata that were to be set up in a Drupal environment.....	13
Figure 4. Development Indico platform showing placeholder training event entries.....	15
Figure 5. Miro canvas showing deliberations and developer notes on the integration between Indico and Drupal (C2ISS), including data to be updated via APIs.....	16
Figure 6. Indico API settings that can be configured with authentication tokens.	17
Figure 7. API access to HPC skills tree information via APIs as documented by the HPC Certification Forum.	19
Figure 8. The concept of SSO and identity providers for access to web-based services and applications.	20
Figure 9. Typical SSO process using SAML based on the roles of user, identity provider and service provider.	21
Figure 10. Miro canvas showing conceptual user registration processes that takes into account SSO for access to other services.....	22
Figure 11. Zenodo community to store training materials uploaded via C2ISS.	23
Figure 12. Initial drafts (on a Miro canvas) of pages to display lists of materials/events as well as individual entries. The visual theme here is outdated and has since been changed entirely.....	24
Figure 13. Miro canvas showing conceptual drafts of features, outlines and designs of an "Attendees' Corner".....	25
Figure 14. First preview of the revamped HPC in Europe Portal landing page.....	26
Figure 15. Previews of training events pages - list of courses (left) and description of a specific course (right). Text and details are to be populated with real data in early 2025.	27
Figure 16. Previews of training material pages - list of materials (left) and description of a specific material (right). Text and details are to be populated with real data in early 2025.	28

Table of Contents

Document Information	2
Document History	2
Executive Summary	3
List of Abbreviations.....	4
List of Figures.....	5
Table of Contents	6
1 Introduction	8
2 Architecture Definition of the Training Platform	8
3 Main Concepts	10
4 Deployment Activities and Tasks.....	12
4.1 Creation of Content Types on C2ISS	12
4.2 Creation of Taxonomies on C2ISS.....	13
4.3 Development of the Indico Platform	14
4.4 Development of API for Data Exchange and Updates	16
4.5 Automation of Data Exchange and Updates	20
4.6 Establishment of Single Sign-On between Drupal and Indico	20
4.7 Setting Up a Materials Repository (via Zenodo).....	23
4.8 Creation of Content Pages (Listing & Details).....	24
4.9 Setting up the Attendee Corner and User Dashboard.....	25
5 Preview of Training Platform Pages	26
5.1 Listing of Training Events & Details.....	27
5.2 Listing of Training Materials & Details.....	28
6 Roll-out Plan of the Pilot Training Platform	29
7 References	30

PUBLIC DELIVERABLE

©2024-25 MEMBERS OF THE HPC SPECTRA CONSORTIUM



1 Introduction

The HPC SPECTRA project addresses two main areas of priority in the EuroHPC JU's training strategy. Firstly, it aims to develop an EuroHPC Training Platform that promotes and advances European HPC training. Secondly, it collaborates with international partners to organise the 2024 and 2025 editions of the International HPC Summer School. This deliverable describes the deployment phase of the training platform for the first objective of the project.

The training platform is intended to serve as a central hub for HPC related training information and access point to training initiatives across Europe. It caters for European training providers to promote training offers (e.g. courses, workshops) and facilitates platform users to identify and find relevant courses and opportunities. The platform will also support the development of a competency framework to better elucidate the competencies being covered by the training offerings in a standardised manner and to help establish accreditation and certification mechanisms. The platform, centred on development of HPC competencies, will be a key pillar in the European HPC ecosystem.

This deliverable re-caps some of the key architectural concepts and considerations that were conceived earlier in this project (Section 2) as well as providing an overview of the main elements of the training platform (Section 3). This is followed by an overall description of various deployment tasks (Section 4) that were implemented in the project, from creation of content types and taxonomies on Drupal, the deployment of external event management platforms using Indico to establishing ways to ensure smooth user management and exchange of information between the different elements. Preview screenshots of the training platform are included in Section 5, which are expected to be populated with live data in January 2025. Finally, Section 6 describes the roll-out plans for the training platform in the second half of the project.

2 Architecture Definition of the Training Platform

Prior to deployment efforts to realise the EuroHPC Training Platform, the HPC SPECTRA project set out to define the architecture of the platform, i.e. the main elements or components of the platform, how they relate and interact with each other,

etc. [1] It had already been established that HPC SPECTRA would not build a new standalone training platform but rather join efforts with the CASTIEL 2 project [2] to revamp the HPC in Europe portal [3] that had intended to serve as a comprehensive directory to European HPC activities, including but not limited to training. This meant that HPC SPECTRA will be responsible for developing the training section of the new, comprehensive pan-European portal, entitled C2ISS [4]. The overall structure of the C2ISS portal is shown in Figure 1, where the training section is to be developed by the HPC SPECTRA project with the incorporation of some external components outside of Drupal (as outlined in Figure 2).

Given the initial conditions of developing the training section of an over-arching HPC information/services portal, the project has outlined and described some of the architectural requirements that are needed to address the main target audience and provide key functionalities [1]. Apart from elucidating the datasets and metadata that will be required, D2.1 proposed the main components that will form the new training platform, including C2ISS (a website built on the Drupal content management system or CMS), integration of a popular event management platform (Indico) and leveraging the capability of Zenodo for longer-term archival of training materials. How these elements combine together and example user workflows have also been explored to ascertain the overall architecture of the platform, which were to be implemented in the deployment phase and described here.

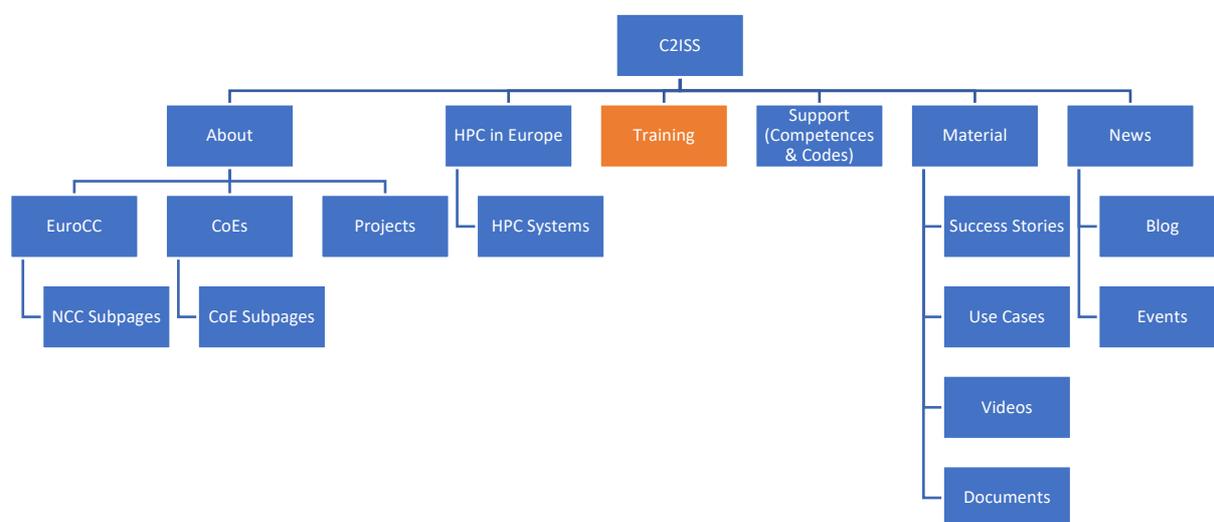


Figure 1. High-level overview of the C2ISS portal structure. Development of the Training section (highlighted in orange) is the responsibility of the HPC SPECTRA project.

3 Main Concepts

The interplay between C2ISS and other training-related components/services are indicated in Figure 2. It shows that while the main training portal will be integrated into the Drupal-based C2ISS, it will interact with external servers (e.g. Indico for event management) and other web services (e.g. Zenodo for materials repository, HPC Certification Forum for skills tree information).

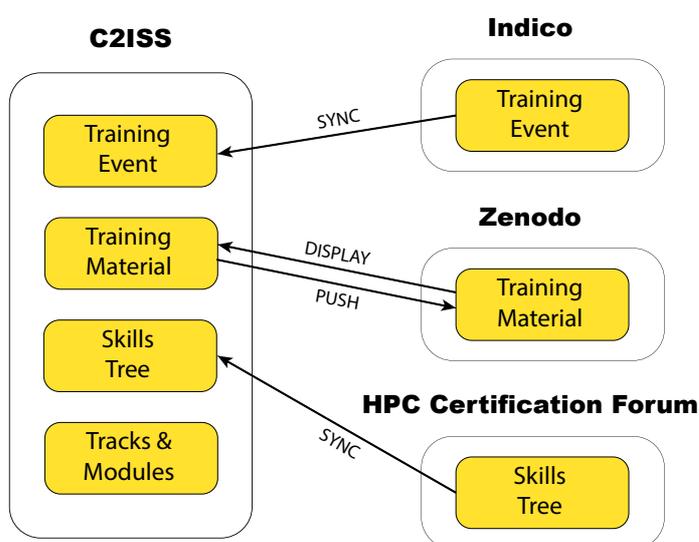


Figure 2. Interplay between C2ISS portal and other training components/services.

Some of the key concepts that are to be implemented in the deployment phase are:

1. **Content and Categorisation**

Training metadata is structured into distinct types and taxonomies to enable robust management and facilitate searches/filtering by users, as well as for ecosystem or landscape analyses (e.g. abundance/paucity of courses in particular areas). The metadata had been defined to strike a balance between usefulness (i.e. whether it enables better filtering, analyses) and burden on the training providers.

2. **Seamless Integration**

The deployment phase involved setting up the necessary APIs and SSO mechanisms to ensure smooth interaction between C2ISS and other components such as Indico and Zenodo. For training providers, for example,

this means the same set of login credentials allows one to add content to C2ISS (Drupal-based) as well as on Indico (where training event metadata is entered).

3. **User Experience**

Realisation of the training platform also involved the drafting of workflows and “user journey” concepts to ensure an intuitive interface for three primary user roles: training providers, portal users (potential attendees) and content managers. Features such as dashboards and user-specific functionalities have been implemented to enhance overall usability.

4. **Automation**

Automated data synchronisation mechanisms (e.g. cron jobs) have been implemented to ensure the different components remains up to date with minimal manual intervention.

4 Deployment Activities and Tasks

The main deployment activities were carried out in WP3, led by KIFU. Regular meetings every two weeks or so were held to coordinate various tasks and actions. The team made extensive use of the Miro collaboration platform [5], an online freeform canvas tool that facilitates collaboration by different groups/individuals, to draw up key design concepts, workflows, content types, interactions, etc. The following subsections (4.1 to 4.9) describe the main deployment tasks carried out by the team.

4.1 Creation of Content Types on C2ISS

Based on the datasets and metadata that have been defined previously, the deployment phase involved setting up the corresponding data fields in a Drupal environment.

- **Training Content:** set of metadata that describes training courses/events from key details (e.g. date, time, mode) to relevant domains and skills covered. Organisers and co-organisers required flexibility not only to cater for different providers from NCCs, CoEs to potentially other European projects and individual institutions; these are maintained as lists which can be expanded upon registration of new training provider accounts.
- **Training Material:** these describe resources such as documents, presentations, and recordings from training events, as well as standalone materials that are uploaded by training providers.
- **Skills:** these are derived from the HPC Certification Forum, the lead of which is also a partner in the project. The skills attributes were set up on Drupal but can be updated via APIs (see below).
- **Element/Module:** these are simply groupings or sets of skills that can become elements of learning pathways or tracks.
- **Track:** Organised sequences of training modules focused on specific learning objectives or technical/career profiles. These tracks represent learning pathways that serve to guide users to relevant courses/materials.

Figure 3 shows a high-level snapshot of the Miro canvas that had been deliberated among WP3 to guide and clarify the implementation of different data types and requirements on the C2ISS Drupal platform. It contains notes on issues to be resolved

for some of the fields, or to emphasise potential accommodation for changes. It also includes the first conceptual sketches (i.e. but not the final visual theme) of the different datasets on a Drupal platform.



Figure 3. Miro canvas showing conceptual workings and sketches of different sets of metadata that were to be set up in a Drupal environment.

4.2 Creation of Taxonomies on C2ISS

Drupal has built-in support for taxonomies that can be used to classify different content types. A simple example is a defined set of keywords that can be used to annotate events or materials, where the keywords conform to a defined hierarchy, e.g. a tree-like or flat structure. Depending on the hierarchy, the taxonomy can easily be expanded on-the-fly according to user input. It is important to note that employment of the taxonomy concept allows for a defined but dynamic set of keywords that are useful in many cases for annotation of training content, i.e. the taxonomies are not set in stone but can evolve over time according to further feedback and requirements. Some of the metadata fields where taxonomies have been employed for training content are as follows:

- Scientific domains (e.g. fundamental physics, not scientific domain specific)
- Skill levels (Beginner, Intermediate, Advanced)
- Target audience (e.g. data scientists, HPC application users)

- Organiser and co-organiser: the list of training providers was to be deployed as a taxonomy that can be expanded dynamically, rather than freeform text, to avoid variations of the same organisation/project name being used, e.g. BSC vs Barcelona Supercomputing Center/Centre. This may be important consequently if one were to look for all BSC courses in future or retrospectively.

4.3 Development of the Indico Platform

As elaborated in the architectural design of the training platform, the HPC SPECTRA project aimed to establish a dedicated Indico platform for training providers. Indico is a popular and open source event management platform that was developed by CERN [6] and has been used extensively by HPC training providers across Europe (e.g. previously deployed for PRACE courses for nearly a decade, a few NCCs have also set up local Indico instances to manage events/courses) as well as being used by the United Nations [7]. It is well maintained and provides the required functions and flexibility for the purposes of the EuroHPC training platform. Rather than just storing events metadata, it includes built-in processes for managing events such as registration, selection/review procedures (if necessary), scheduling, feedback collection, etc. Events on Indico are also semantically interpretable by search engines, i.e. Google searches can return results as events (with details such as date, time, location) rather than simple links to web pages.

During the deployment phase, some of the sub-tasks involved in the deployment of Indico include the following:

- Several Indico instances (test / staging / production) were set up on virtual machines to test different functionalities.
- For training events, customisation of events metadata was carried out to ensure consistency with those that were derived for training courses/events.
- SSO integration with Drupal was investigated by examining compatible authentication methods and any additional modules that may be required to be installed on either Indico/Drupal (see subsection 4.6 below).
- Different user roles were set up to be compatible with account registrations that were to be implemented on the C2ISS platform, i.e. users who are able to create

training events on Indico. This also involved setting up a developer roles who have administrator access on the Indico platform.

- Categorisation and hierarchy of training events were set up in order to accommodate groups of similar events on Indico (Figure 4).
- Incorporate visual themes that are consistent with the C2ISS platform.

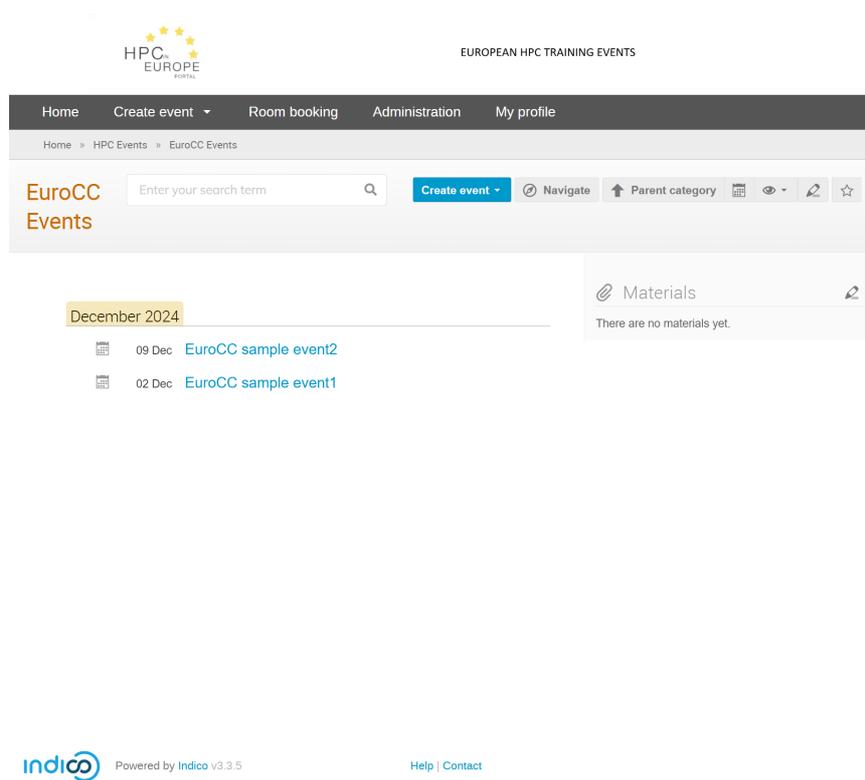


Figure 4. Development Indico platform showing placeholder training event entries.

4.4 Development of API for Data Exchange and Updates

This important task was focused on maintaining consistency of data between the Indico platform and the C2ISS Drupal CMS. As training event information is entered on Indico, the main flow of data is from Indico to the Drupal system. Figure 5 shows the workings, discussions, notes and deliberations on the type of data that needs to be sent and the potential means and formats. One of the main ways to exchange data between these systems is the use of Application Programming Interface (API), which is often supported natively or via popular extensions/modules. All the components of the training platform support the use of APIs to send/receive data. Apart from data exchange between Indico and Drupal, this task was also involved in using APIs to exchange information with the Zenodo archival service (for depositing training materials and metadata) and the skills tree framework (developed by the HPC Certification Forum).

In order to ensure also consistency of users (e.g. training providers) across Drupal and Indico, this is achieved via Single Sign-On implementation which is discussed further in Section 4.6.

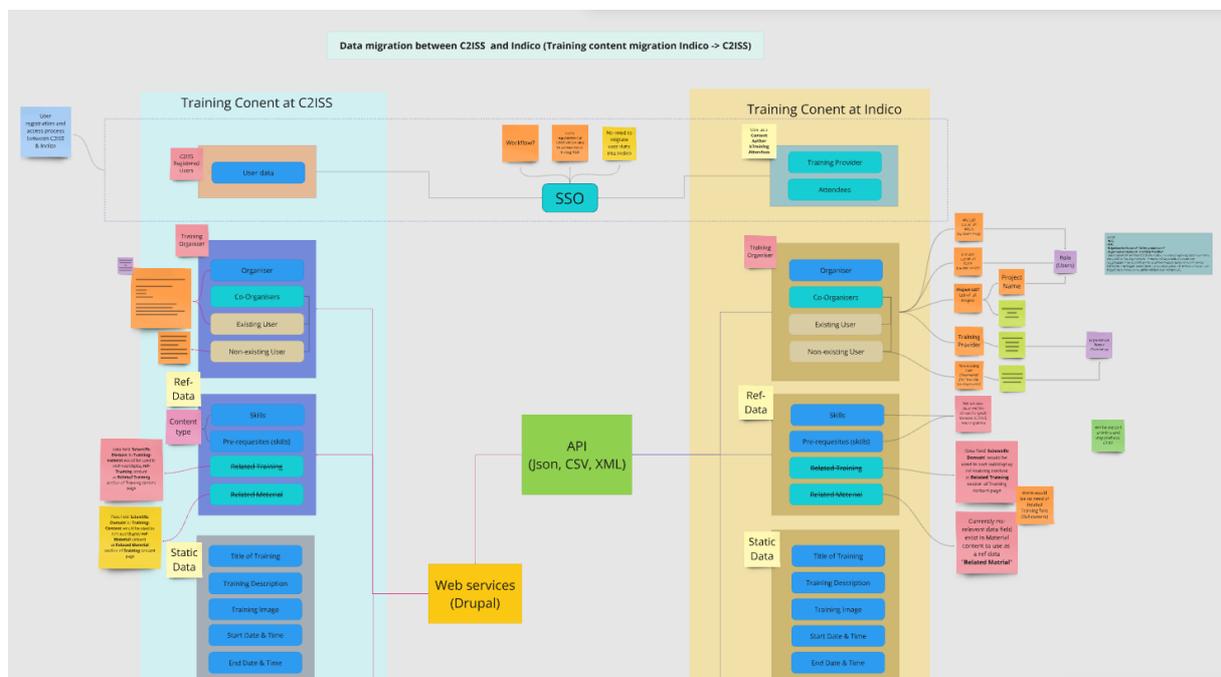
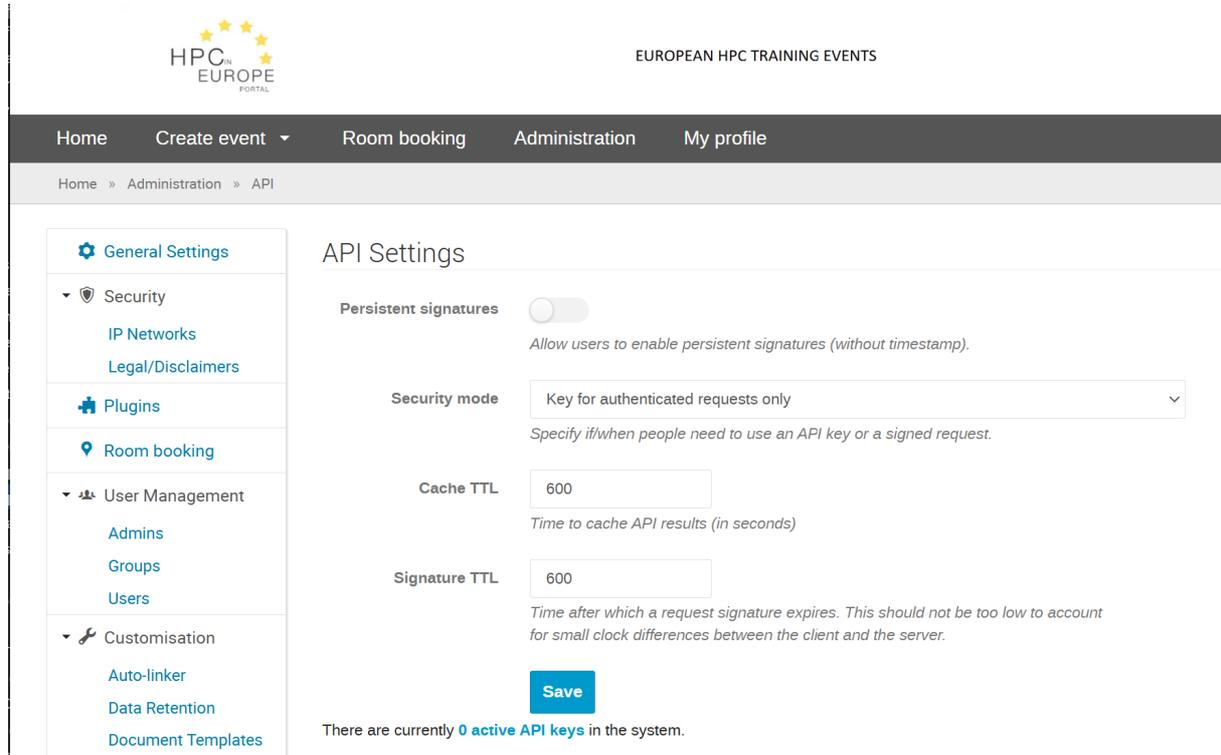


Figure 5. Miro canvas showing deliberations and developer notes on the integration between Indico and Drupal (C2ISS), including data to be updated via APIs.

Integration of Indico and Drupal requires using Indico API [8] to feed data to the Drupal side. This can be achieved via utilising Indico HTTP API. One can create user-managed API tokens to allow limitation of access for specific targets.



The screenshot shows the 'API Settings' page in the HPC in Europe Portal. The page has a navigation bar with 'Home', 'Create event', 'Room booking', 'Administration', and 'My profile'. The breadcrumb trail is 'Home » Administration » API'. The left sidebar contains a menu with 'General Settings', 'Security' (with sub-items 'IP Networks' and 'Legal/Disclaimers'), 'Plugins', 'Room booking', 'User Management' (with sub-items 'Admins', 'Groups', and 'Users'), and 'Customisation' (with sub-items 'Auto-linker', 'Data Retention', and 'Document Templates'). The main content area is titled 'API Settings' and contains the following settings:

- Persistent signatures:** A toggle switch is currently turned off. Below it, the text reads: 'Allow users to enable persistent signatures (without timestamp).'.
- Security mode:** A dropdown menu is set to 'Key for authenticated requests only'. Below it, the text reads: 'Specify if/when people need to use an API key or a signed request.'
- Cache TTL:** A text input field contains the value '600'. Below it, the text reads: 'Time to cache API results (in seconds)'.
- Signature TTL:** A text input field contains the value '600'. Below it, the text reads: 'Time after which a request signature expires. This should not be too low to account for small clock differences between the client and the server.'

At the bottom of the settings area, there is a blue 'Save' button and a status message: 'There are currently 0 active API keys in the system.'

Figure 6. Indico API settings that can be configured with authentication tokens.

The API works on URL paths similar to the following:

```
/export/event/137346.json?occ=yes&pretty=yes
```

The above is an example of a JSON (JavaScript Object Notation) API access path. JSON is a commonly used and flexible format used in data exchange between web services.

The following demonstrate an example of an API response to a call for metadata about a specific event:

```
{
  "count": 1,
  "additionalInfo": {

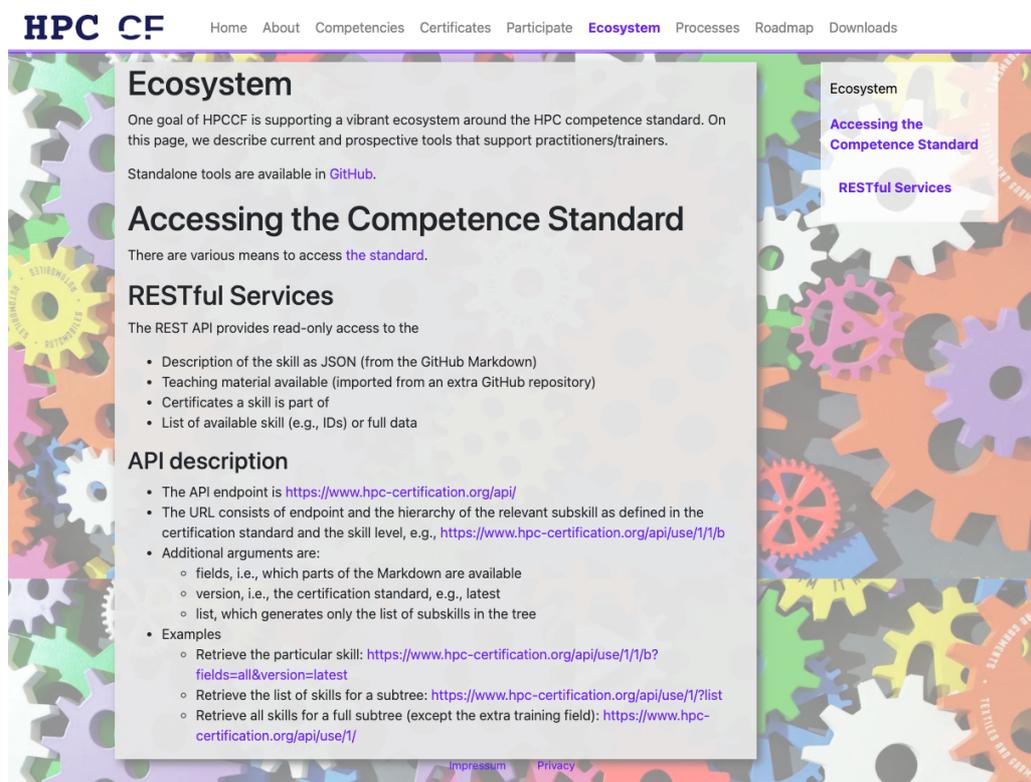
  },
  "ts": 1734604190,
  "url": "https://spectra-indico.hpc.einfra.hu/export/event/2.json?occ=yes&pretty=yes",
  "results": [
    {
      "_type": "Conference",
```

```
"id": "2",
"title": "EuroCC sample event2",
"description": "",
"startDate": {
  "date": "2024-12-09",
  "time": "16:00:00",
  "tz": "Europe/Budapest"
},
"timezone": "Europe/Budapest",
"endDate": {
  "date": "2024-12-09",
  "time": "18:00:00",
  "tz": "Europe/Budapest"
},
"room": "",
"location": "",
"address": "",
"type": "conference",
"references": [],
"_fossil": "conferenceMetadata",
"categoryId": 2,
"category": "EuroCC Events",
"note": {
},
"roomFullname": "",
"url": "https://spectra-indico.hpc.einfra.hu/event/2/",
"creationDate": {
  "date": "2024-12-02",
  "time": "15:29:34.448531",
  "tz": "Europe/Budapest"
},
"creator": {
  "_type": "Avatar",
  "_fossil": "conferenceChairMetadata",
  "first_name": "Zoltán",
  "last_name": "Kiss",
  "fullName": "Kiss, Zoltán",
  "id": "5",
  "affiliation": "employee@niif.hu;staff@niif.hu;member@niif.hu",
  "emailHash": "577152380e07c00b485c4fd3ba0a1811"
},
"hasAnyProtection": false,
"roomMapURL": "",
"folders": [],
"chairs": [],
"material": [],
"keywords": [],
"organizer": "",
"language": null,
"label": null,
"visibility": {
  "id": "",
  "name": "Everywhere"
},
"occurrences": []
}
],
"_type": "HTTPAPIResult"
}
```

The C2ISS Drupal system makes API requests to Indico in order to create or update training events, which could in turn be displayed, indexed and searched via filtering options. As mentioned above, C2ISS has been readily configured with the corresponding data types and structures in order to store such information.

The use of APIs to deposit training materials and metadata to Zenodo for longer-term archival is discussed further in Section 4.7. Again, C2ISS remains to be the main display and gateway to the materials (i.e. where user may browse or search for relevant materials), but the materials themselves along with the metadata resides within the Zenodo system, allowing for such materials to be also accessible elsewhere via other mechanisms (e.g. web searches) and stored with unique, permanent Document Object Identifiers.

On the C2ISS platform, the skills tree data can also be populated and updated via API calls to web services set up by the HPC Certification Forum, whose lead is involved in the HPC SPECTRA project. Figure 7 shows how the relevant documentation on API access to the skills tree information.



The screenshot shows the HPC CF website with a navigation menu: Home, About, Competencies, Certificates, Participate, **Ecosystem**, Processes, Roadmap, Downloads. The main content area is titled "Ecosystem" and contains the following sections:

- Ecosystem**: One goal of HPCCF is supporting a vibrant ecosystem around the HPC competence standard. On this page, we describe current and prospective tools that support practitioners/trainers. Standalone tools are available in [GitHub](#).
- Accessing the Competence Standard**: There are various means to access [the standard](#).
- RESTful Services**: The REST API provides read-only access to the
 - Description of the skill as JSON (from the GitHub Markdown)
 - Teaching material available (imported from an extra GitHub repository)
 - Certificates a skill is part of
 - List of available skill (e.g., IDs) or full data
- API description**
 - The API endpoint is <https://www.hpc-certification.org/api/>
 - The URL consists of endpoint and the hierarchy of the relevant subskill as defined in the certification standard and the skill level, e.g., <https://www.hpc-certification.org/api/use/1/1/b>
 - Additional arguments are:
 - fields, i.e., which parts of the Markdown are available
 - version, i.e., the certification standard, e.g., latest
 - list, which generates only the list of subskills in the tree
 - Examples
 - Retrieve the particular skill: <https://www.hpc-certification.org/api/use/1/1/b?fields=all&version=latest>
 - Retrieve the list of skills for a subtree: <https://www.hpc-certification.org/api/use/1/?list>
 - Retrieve all skills for a full subtree (except the extra training field): <https://www.hpc-certification.org/api/use/1/>

At the bottom of the page, there are links for [Impressum](#) and [Privacy](#).

Figure 7. API access to HPC skills tree information via APIs as documented by the HPC Certification Forum.

4.5 Automation of Data Exchange and Updates

The synchronisation of data between the different systems using APIs can be ensured by scheduling periodic update requests on the relevant servers, e.g. the use of the Linux cron jobs and shell scripts. The frequency of these updates varies depending on the nature of the data. So while the skills tree information is unlikely to require frequent updates, training events information should be reflected on relevant systems within minutes; these frequencies can be refined as more users are onboarded onto these systems.

4.6 Establishment of Single Sign-On between Drupal and Indico

The concept of Single Sign-on (SSO) has been mainstream on the web that saves users from having to set up multiple accounts on different websites. Examples include the use of commonly used Google, Microsoft credentials in order to access other websites and services. In these, Google and Microsoft represent the Identity Providers (IdPs) that facilitates SSO functionality. Apart from large corporations, many databases and open source technologies could also be adopted as identity providers in other settings (Figure 8).

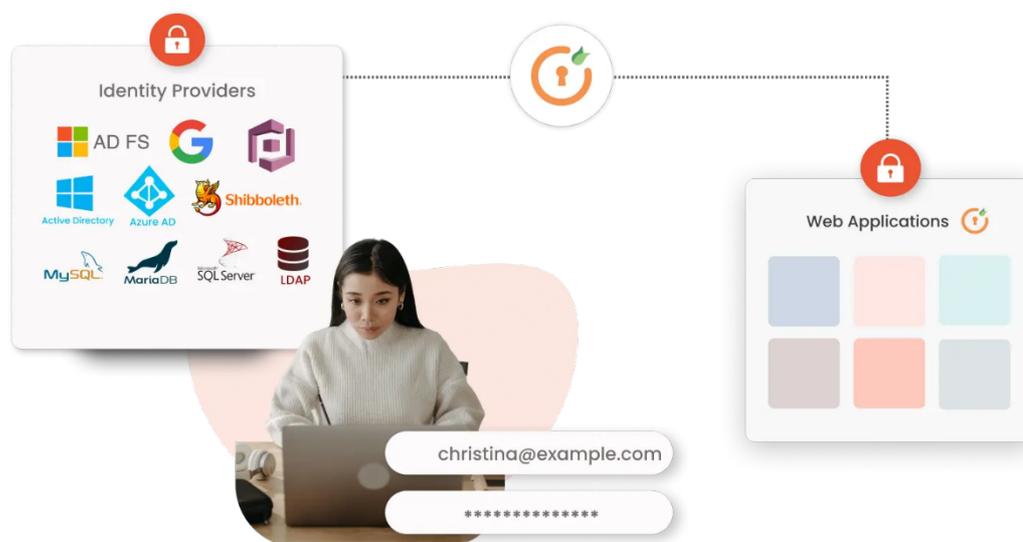


Figure 8. The concept of SSO and identity providers for access to web-based services and applications.

For the training platform, Drupal was established as the IdP to allow the Indico event management system to use the same information for handling user logins and consequently access privileges and data. A Drupal add-on called miniOrange [9] was deployed to serve as an identity provider through SAML (Security Assertion Markup Language). SAML is an open standard that allows IdPs to exchange or pass on authentication and authorisation data between different service providers. It offers security features to protect unauthorised access and data privacy, as well as one of the more commonly used and convenient method of user data exchange.

The SAML specification is primarily based on three roles: the principal (typically a human user), the identity provider (IdP), the service provider (Figure 9). For the training platform, the miniOrange set-up on C2ISS acts as the IdP while Indico represents a service provider, which can obtain all the necessary information to authenticate and handle users.

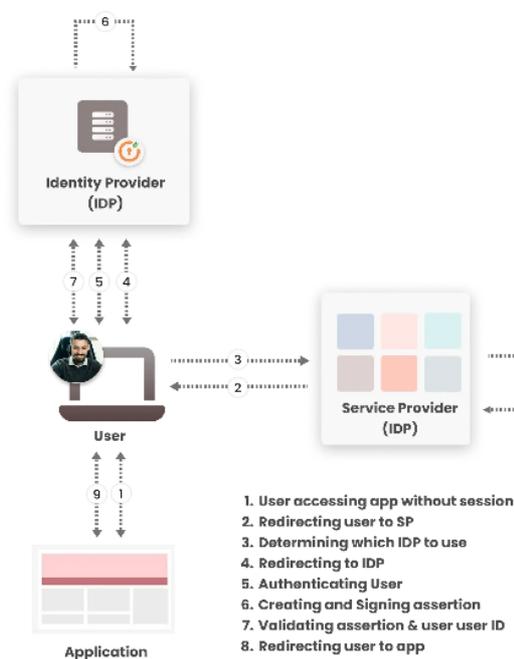


Figure 9. Typical SSO process using SAML based on the roles of user, identity provider and service provider.

In order to properly handle user authentication, roles and privileges, attribute mapping between C2ISS and Indico was carried out to ensure consistency and properly assigned user roles. In addition, there have been much deliberations (Figure 10) on the user registration processes for different types of users, since C2ISS as a platform caters for

not only training providers but the European HPC community including NCC, CoE and European project users as well as potential training attendees (who wish to register an account). These have all been set up in a way to ensure SSO functionality so that an individual would only require a single account that would enable the relevant access to other back-end platforms where necessary.



Figure 10. Miro canvas showing conceptual user registration processes that takes into account SSO for access to other services.

4.7 Setting Up a Materials Repository (via Zenodo)

Apart from training events, the training portal was also being developed to serve as a repository of HPC and related training materials, both from training events as well as standalone materials contributed by training providers. Such a repository had been anticipated to manage a relatively large volume and variety of data (i.e. from documents, presentation slides and notes, video/audio recordings) as well as critical metadata (e.g. author, affiliation, description). As indicated previously, the HPC SPECTRA project had proposed to use Zenodo [10] as a longer-term repository of such materials, as opposed to setting up a dedicated infrastructure that will require more resource-intensive long-term maintenance beyond the life of this project. Furthermore, Zenodo offers a convenient solution (e.g. generation of DOIs) as well as an alternative route to make the datasets of training material more accessible beyond C2ISS and the training portal.

An European HPC knowledge base community on Zenodo had been created allow content to be uploaded and managed (Figure 11). The knowledge base handles all types of materials including images, videos, documents, etc, and supports DOI generation upon upload.

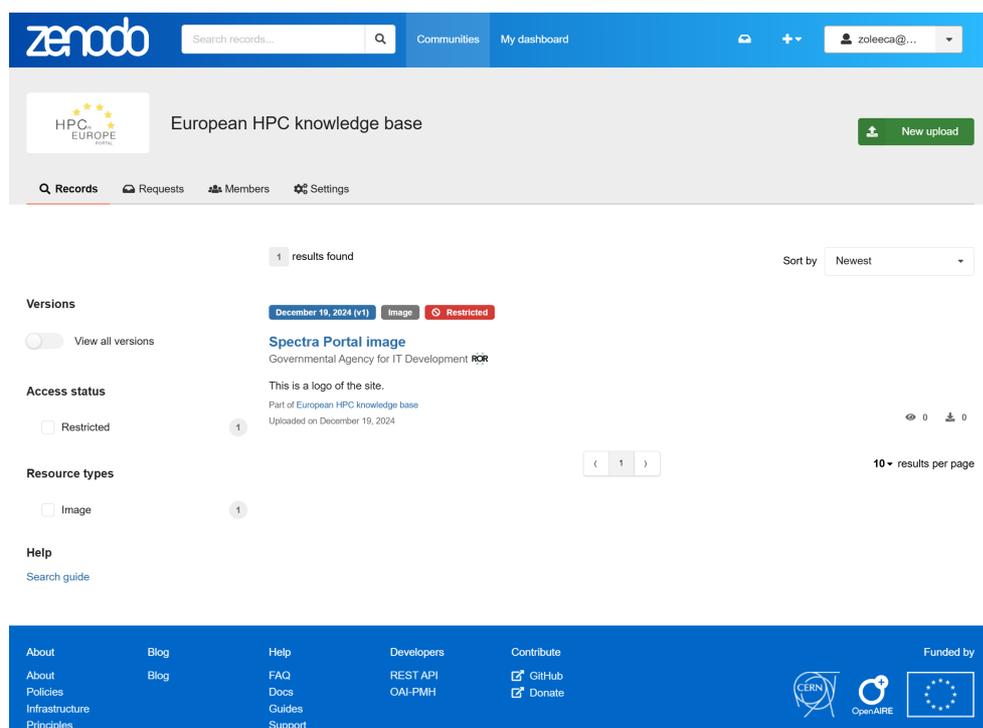


Figure 11. Zenodo community to store training materials uploaded via C2ISS.

As described in Figure 2 and subsection 4.4, Zenodo is integrated with C2ISS (Drupal based system) by implementation of API push and API listing function. A training provider user on the C2ISS platform can specify files to be uploaded along with metadata that are relevant, these are then “pushed” to Zenodo with all the metadata. Conversely, C2ISS remains to be the gateway for the same materials, allowing website visitors to browse, filter and search through the same materials which can be linked to and downloaded seamlessly from Zenodo.

4.8 Creation of Content Pages (Listing & Details)

This sub-task involved developing the conceptual and visual implementation of the pages that forms the main landing pages that displays the list of training events and materials Figure 12. It encompasses not only the overall layout but also interfaces that would allow users to conveniently browse, filter, search for relevant events/materials.

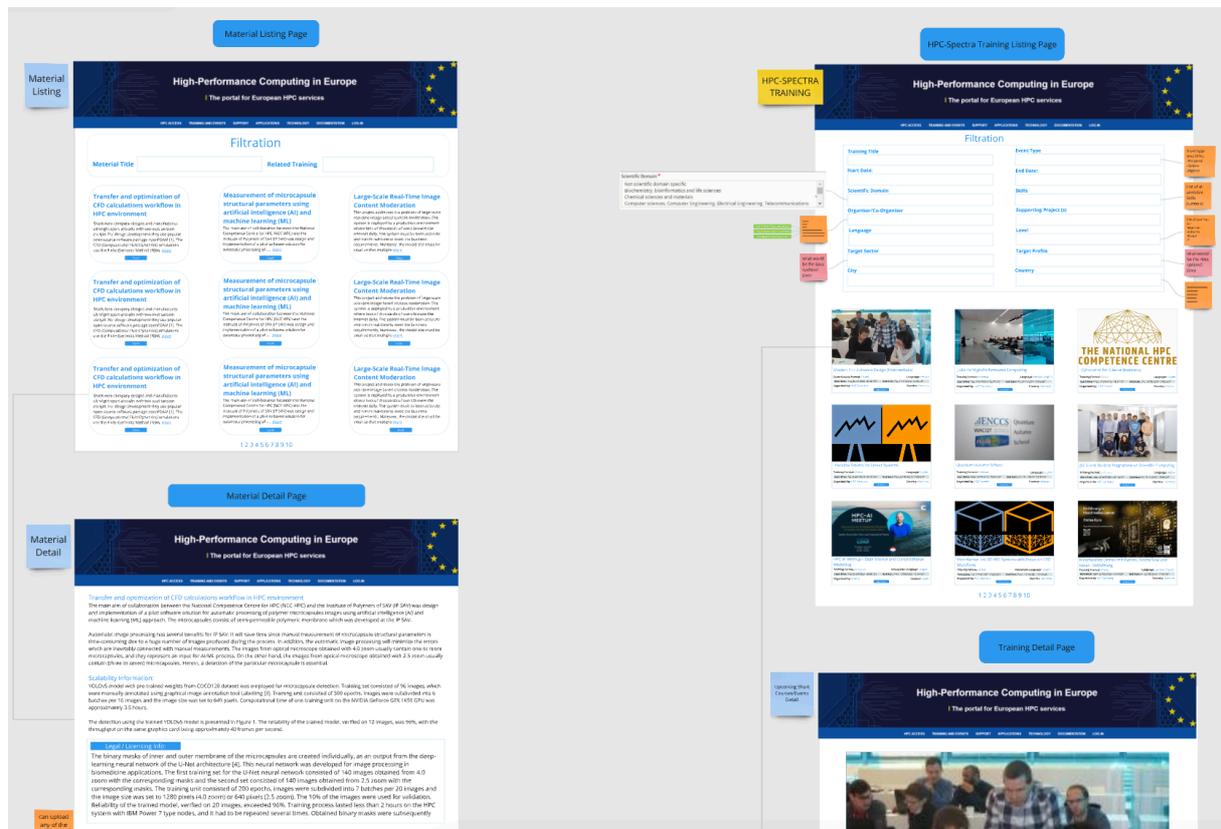


Figure 12. Initial drafts (on a Miro canvas) of pages to display lists of materials/events as well as individual entries. The visual theme here is outdated and has since been changed entirely.

4.9 Setting up the Attendee Corner and User Dashboard

One of the more novel feature of the training platform that had been planned is aimed not at training providers but for trainees (and potential trainees). Here, any individual can register for an account on C2ISS for the following features:

- Provide details on interest in learning about certain skills, or in particular learning pathways based on career/domain profiles. These are all based on skills derived from the skills tree, the “link” to relevant courses.
- Have the system recommend courses (annotated with skills / modules / tracks metadata based on the skills tree) that may be of particular interest to the user based on the profile.
- Keep track of training events that the user had participated in previously.
- A more convenient course registration process for courses that utilises the registration facility of the Indico platform set up by HPC SPECTRA.
- Be directed to relevant mechanisms to attain certification of specific skills, e.g. via assessment platforms to be launched by the HPC Certification Forum.
- Keep track of certifications attained by the user.

Figure 13 shows the conceptual drafts of how such an “Attendee’s Corner” feature were to be implemented on C2ISS. At the time of reporting, further testing of this novel feature is required with additional information about courses before this can be rolled out on C2ISS.



Figure 13. Miro canvas showing conceptual drafts of features, outlines and designs of an “Attendees’ Corner”.

5 Preview of Training Platform Pages

WP3 had worked closely with the CASTIEL 2 project in the implementation and deployment of the training section which coincides with the overall pilot launch of the whole C2ISS or revamped “HPC in Europe” platform, which is due to take place in early January 2025. At the time of reporting, previews of some of the key pages are demonstrated in this section ahead of the first ingestion or population of data on the Drupal CMS. This will also coincide with the induction of new users for initial testing in a pilot phase (see Section 6).

The visual theme and main landing page of the revamped HPC in Europe Portal (i.e. C2ISS) can be previewed in Figure 14, where the training section is prominently indicated in the top menu bar as well as within the main content links.

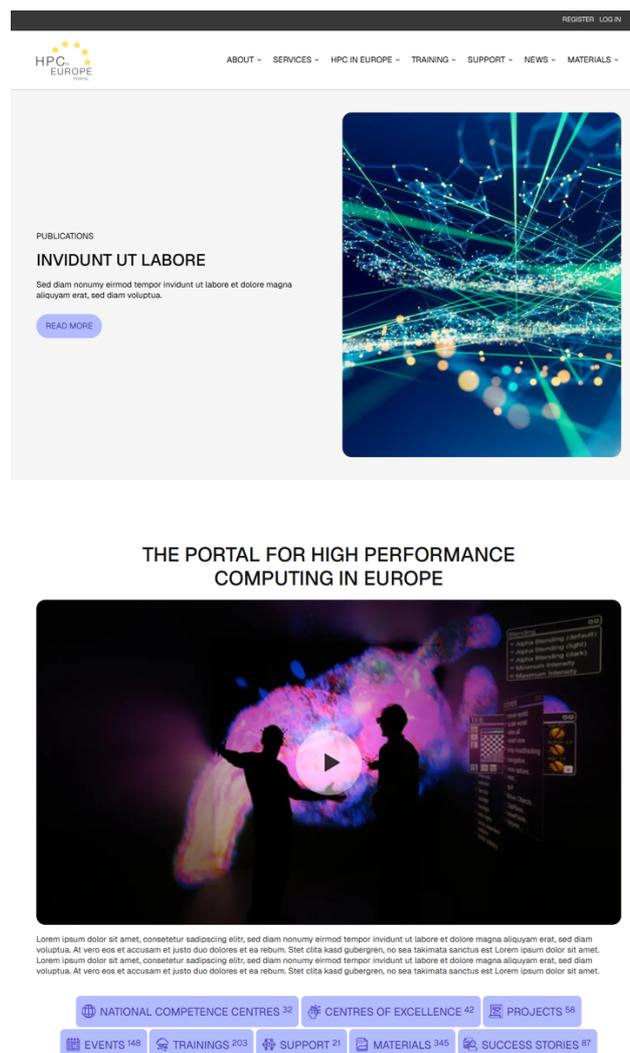


Figure 14. First preview of the revamped HPC in Europe Portal landing page.

5.1 Listing of Training Events & Details

Figure 15 shows screenshot previews of the training section of the HPC in Europ (i.e. C2ISS) platform to (a) display lists of courses, and (b) display information about a specific course. The events listing page includes fields for freetext search as well as filters (e.g. date, location, type) that can be configured by the user to search for specific courses. While the content is filled with placeholder titles and text, these will be updated when data is ingested into the platform in early 2025.

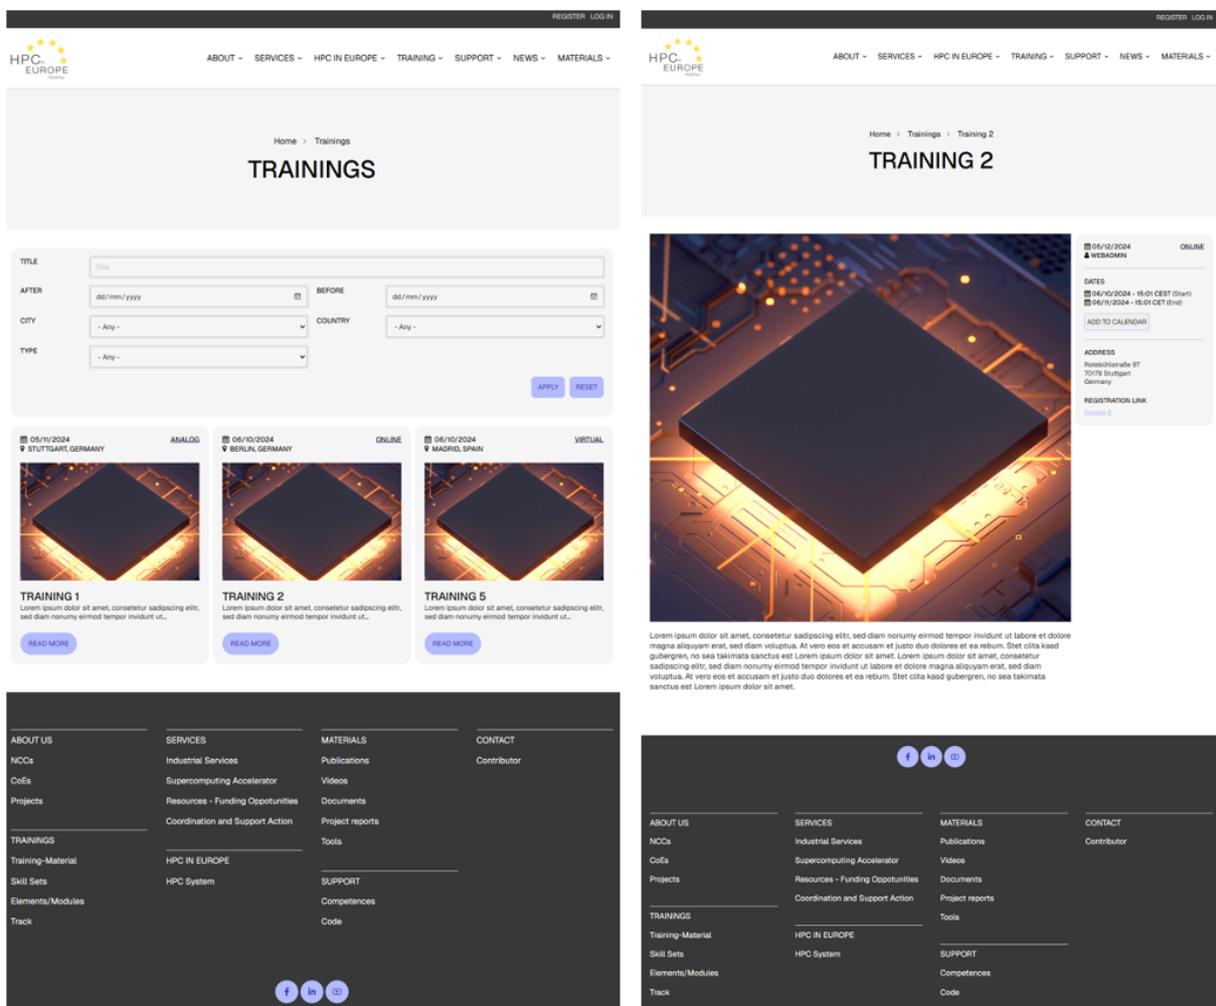


Figure 15. Previews of training events pages - list of courses (left) and description of a specific course (right). Text and details are to be populated with real data in early 2025.

5.2 Listing of Training Materials & Details

Figure 16 shows screenshot previews of the training section of the HPC in Europe (i.e. C2ISS) portal for a list of materials and for a specific material. Again, the list of materials includes filters that can be configured by the user to search for relevant materials. While the content is filled with placeholder titles and text, these will be updated when data is ingested into the platform in early 2025.

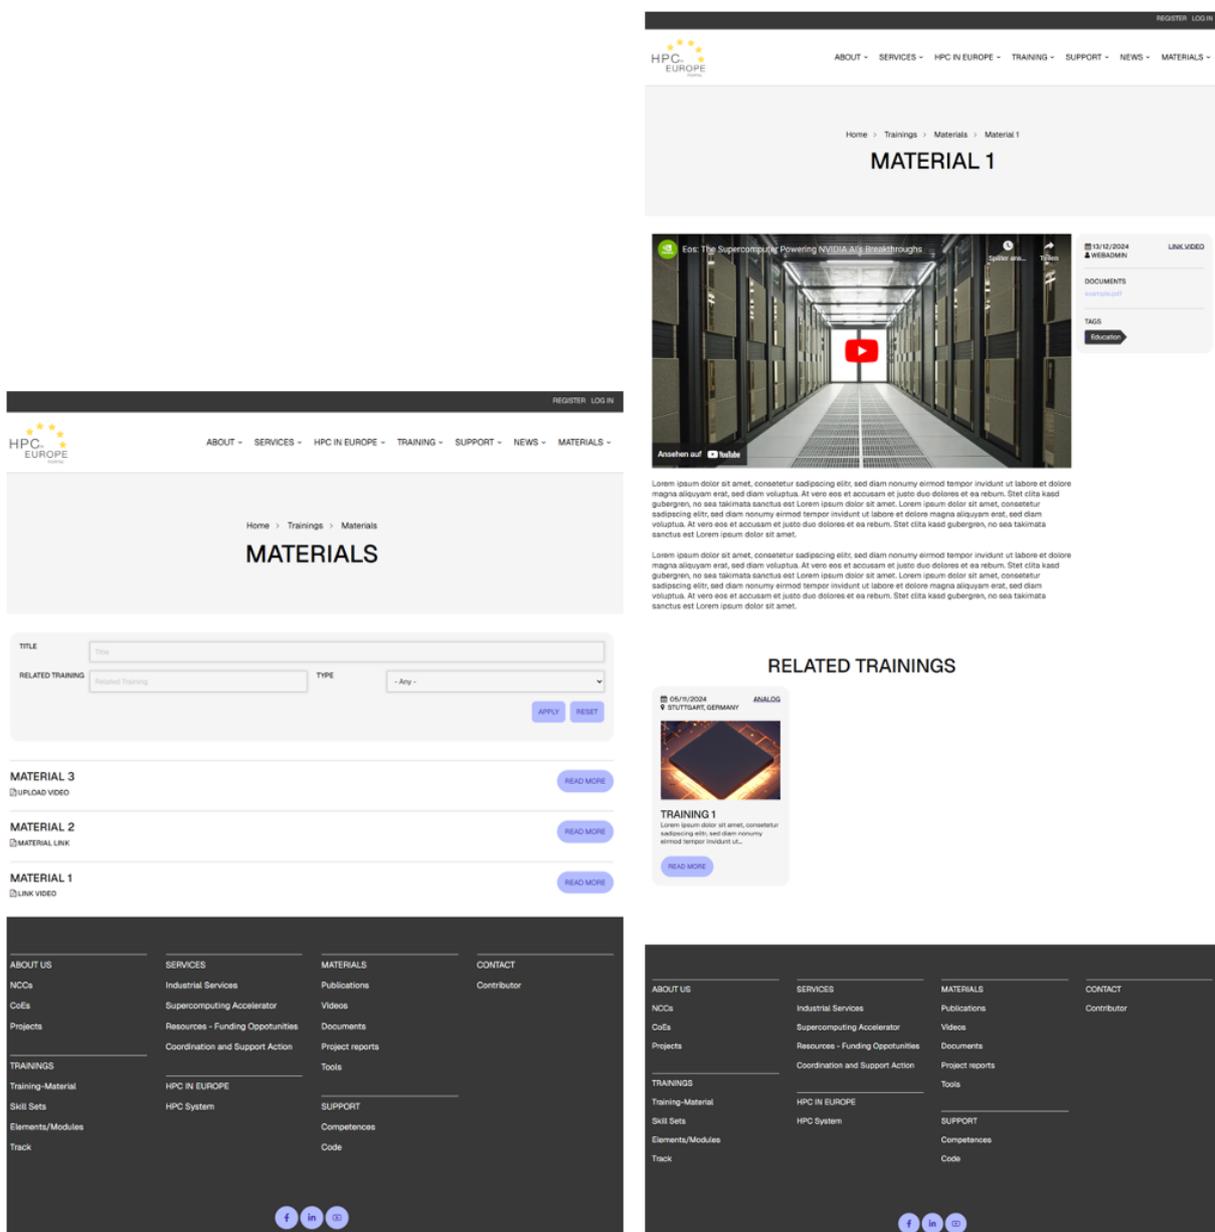


Figure 16. Previews of training material pages - list of materials (left) and description of a specific material (right). Text and details are to be populated with real data in early 2025.

6 Roll-out Plan of the Pilot Training Platform

The HPC in Europe (i.e. C2ISS) Portal had been scheduled to officially launch in January 2025, together with the training components as described in this deliverable. This will coincide with an initial data ingestion period where pre-existing data will be populated on the portal. The roll-out plan for the training platform is as follows:

1. Initial period of limited/graduated user access

It is anticipated that a small number of NCCs will be on-boarded to the training platform as initial adopters in order to obtain initial feedback and identify any significant issues before a broader roll-out to other training providers and users.

2. Further internal testing

Internally, further usability testing will be conducted to ensure a smooth user experience and refine the different workflows and processes in managing the platforms, e.g. potential issues with SSO, user registrations.

3. Feedback mechanism

Ensure that there is a feedback mechanism for portal users to report bugs or highlight issues.

4. Bug fixes and improvements

The deployment team will resolve issues identified during the pilot phase and enhance system features as needed.

5. Dissemination of the platform

There will be an emphasis to continually disseminate the training platform to existing and potential new users, in particular training providers. This will be accompanied by production of onboarding materials for an improved user experience.

It is expected that the roll-out or pilot phase will take place over several months, with the expectation that the training platform will be opened to all users and training providers by mid 2025.

7 References

- [1] HPC SPECTRA project deliverable D2.1: “Architecture of the EuroHPC Training Platform”.
- [2] Information about the CASTIEL 2 project: <https://www.eurocc-access.eu/about-us/the-projects/>
- [3] HPC in Europe Portal: <https://hpc-portal.eu/>
- [4] CASTIEL 2 project deliverable D5.2: “Design strategy of the C2ISS and the evolved EuroCC Access Portal”.
- [5] Miro Collaboration Platform/Workspace: <https://miro.com/>
- [6] Indico software by CERN: <https://indico.cern.ch/>
- [7] Indico deployment at the United Nations: <https://indico.un.org/>
- [8] Indico documentation on its APIs: <https://developer.indicodata.ai/docs/getting-started>
- [9] miniOrange plugin for Drupal: <https://plugins.miniorange.com/drupal>
- [10] Zenodo digital repository: <https://zenodo.org/>