

Updates on InvenioRDM Development at BNL

Collaborators: Dmitry Arkhipkin, Uma Ganapathy, Vincent Garonne, Robert Hancock, Eric Lancon, Jerome Lauret, Louis Pelosi, Ofer Rind, Peter Steinberg

May 1st, 2025

CRISP

Collaborative Research Information Sharing Platform

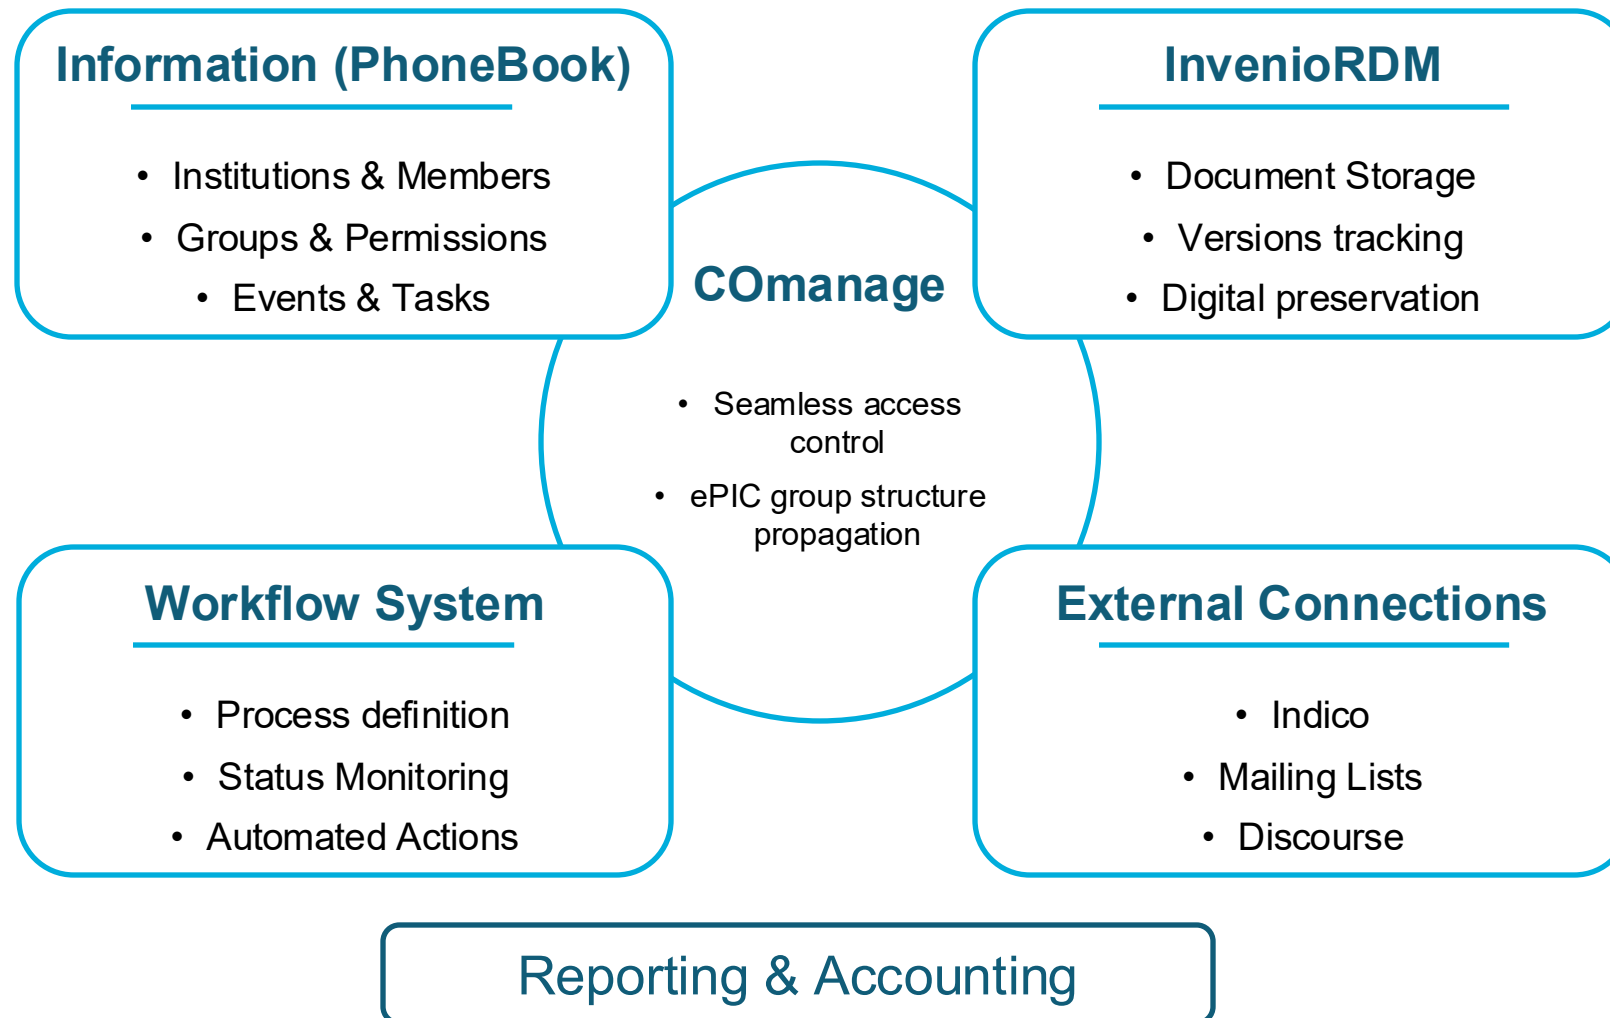
A unified ecosystem enabling seamless scientific collaboration

Capabilities:

- Authentication
- Development
- Communication
- Sharing

Designed by scientists, for scientists with modular and extensible architecture

CRISP: Integrated Platform Components



Need to register with COmanage

- COmanage : One account to access all Web-based services
 - Streamline access using your BNL account or your home institution's credentials.
 - This eliminates the need for non-BNL users to create a BNL-specific account.
 - Accounts from various providers can be combined (BNL, BNL/SDCC, CERN,...) through the use of ORCID.
- COmanage is the tool used to validate your identity against the list of ePIC members provided by the PhoneBook.

COmanage Enrollment Flow

This must be done once.

Detailed instructions are available on the SDCC website

1. Create an Entry in COmanage:

Log in to the SDCC COmanage instance and set up your profile.

2. Phase I - Link to Identity Provider:

Connect your COmanage entry to your institution's Identity Provider (IdP).

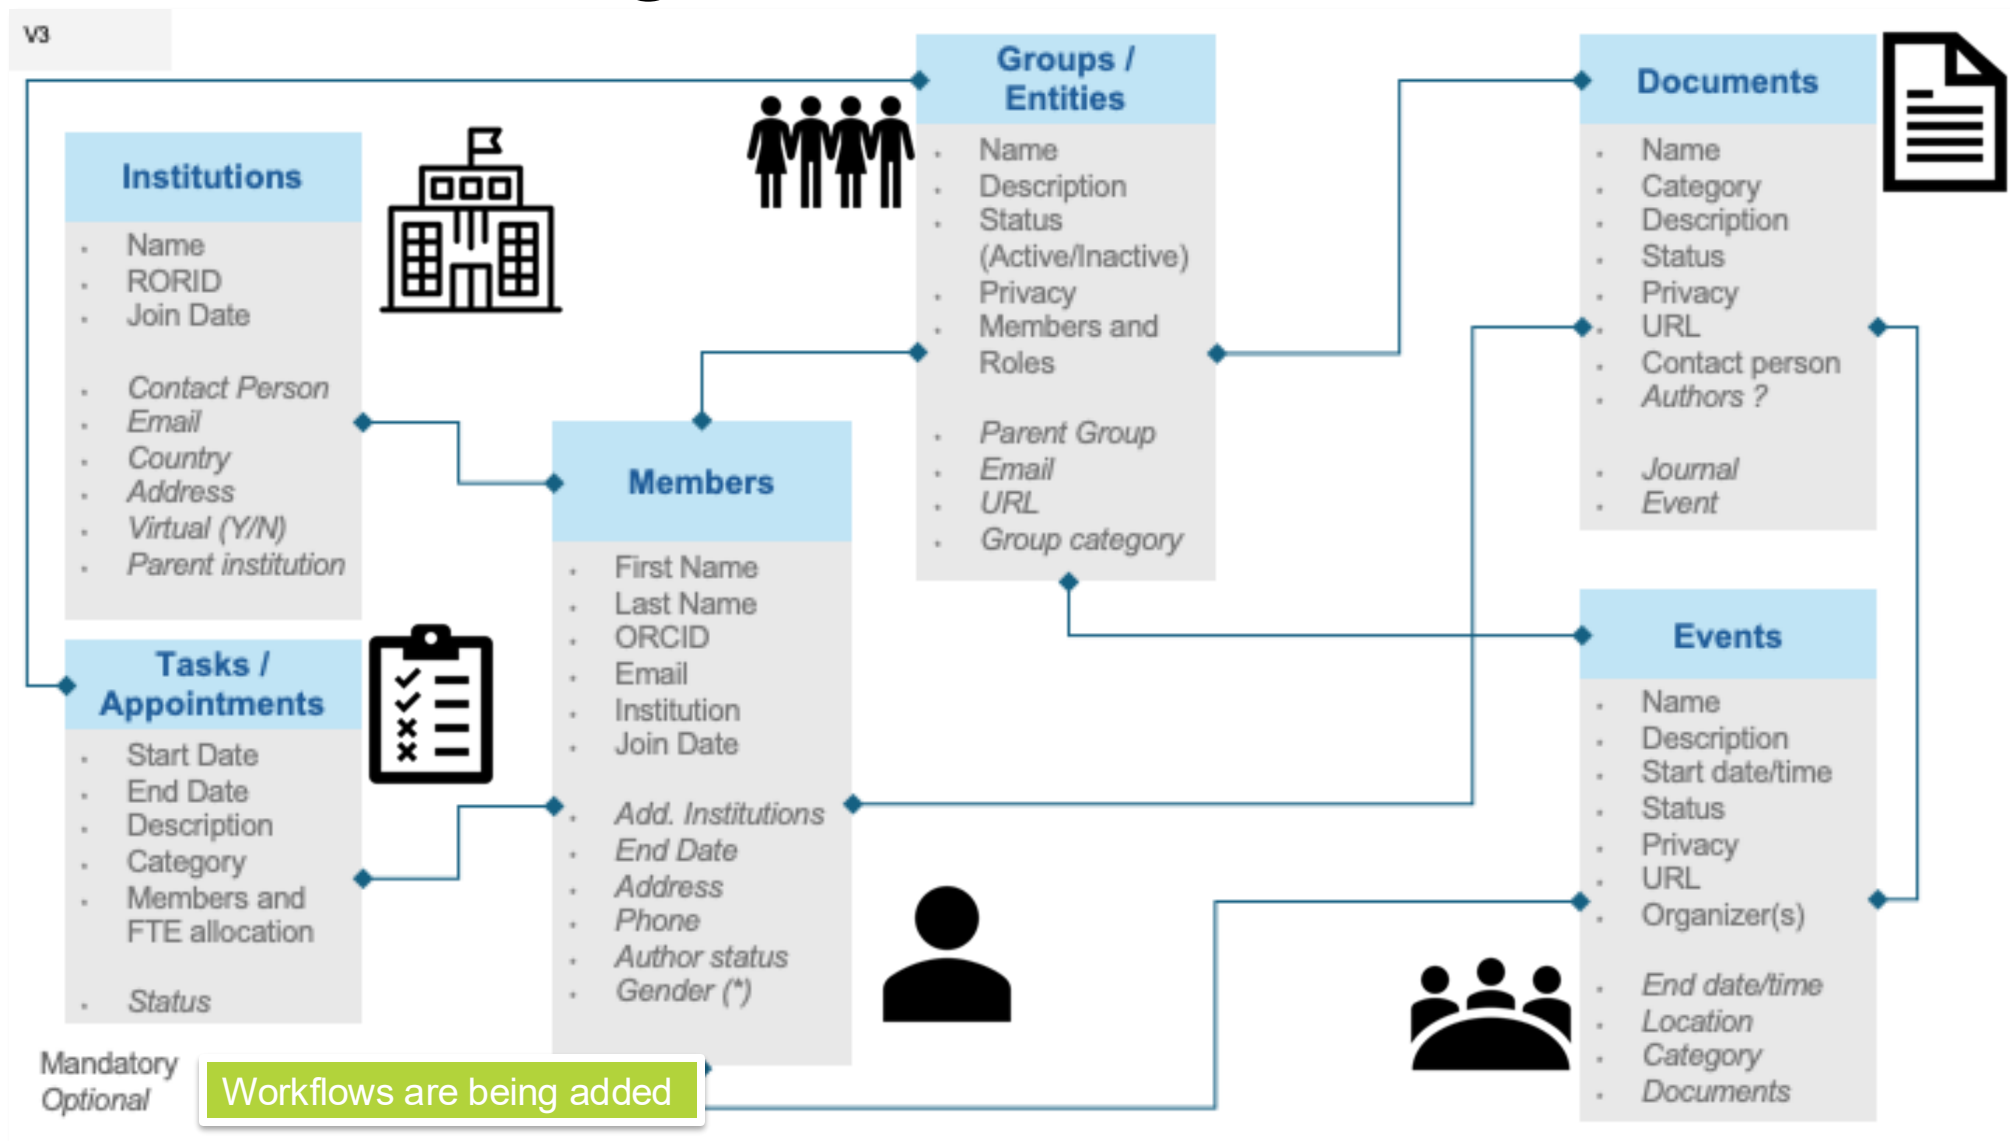
3. Phase II - Email Validation:

Verify your entry by following the link sent to your email.

4. Link to ORCID:

Connect your COmanage entry to your ORCID. This unique identifier helps disambiguate individuals. If you don't have one yet, you can easily register at this [link](#).

PhoneBook Organization



The PhoneBook

Groups:

- Hierarchical layout
- With various roles for members

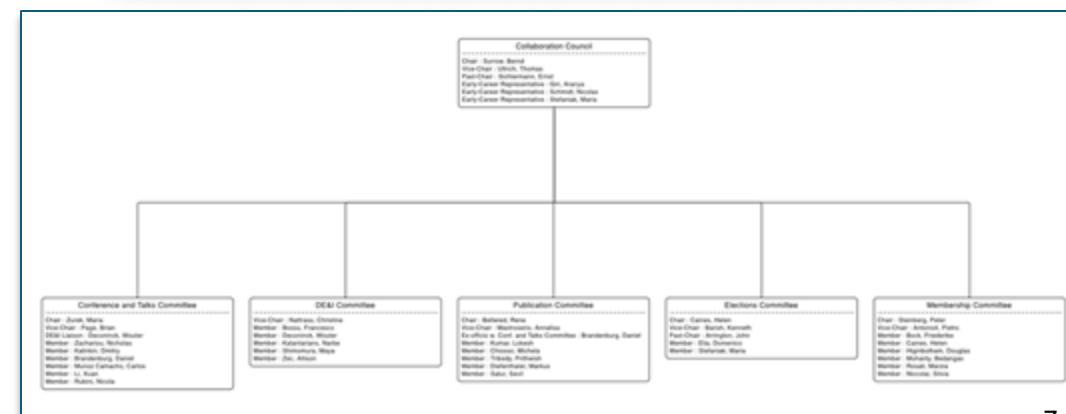
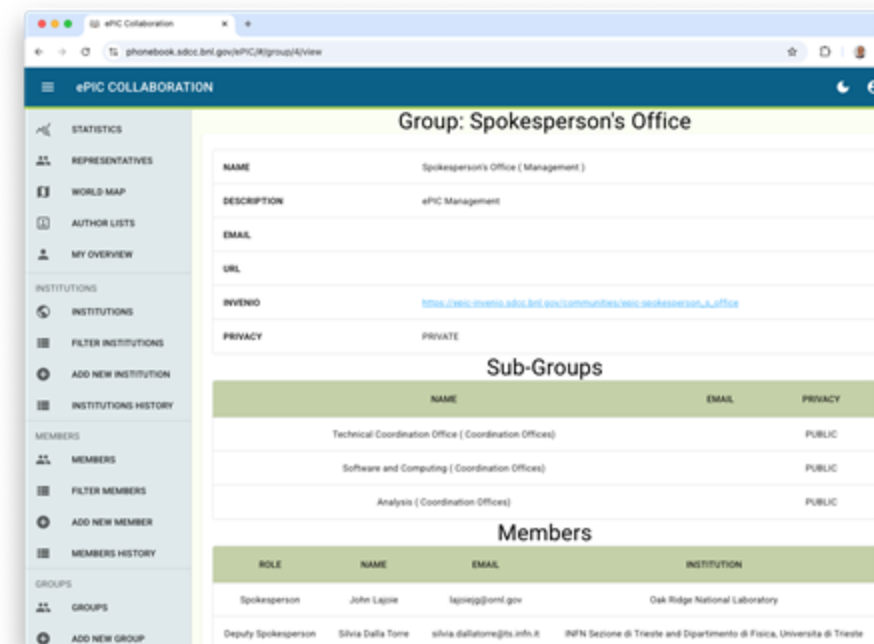
Institutions:

- Defined by their ROR ID
- Can be nested or virtual



Tasks with associated groups,
institutions, and members

Rely on ORCID for member
disambiguation



User view

Personal information

The screenshot shows the 'My Overview' page for a user named Eric Lancon. The page has a dark blue header with the 'ePIC COLLABORATION' logo. A left sidebar contains navigation links: STATISTICS, REPRESENTATIVES, WORLD MAP, AUTHOR LISTS, MY OVERVIEW (highlighted), INSTITUTIONS, FILTER INSTITUTIONS, ADD NEW INSTITUTION, INSTITUTIONS HISTORY, MEMBERS, FILTER MEMBERS, and ADD NEW MEMBER. The main content area has tabs for MY BIO, MY INSTITUTION, and MY GROUPS. Under MY BIO, there is a profile picture of Eric Lancon, his name, email (elancon@bnl.gov), and affiliation (Brookhaven National Laboratory, UNITED STATES). Below this is a table of personal details:

ORCID	0000-0002-0225-187X
FIRST NAME	Eric
FAMILY NAME	Lancon
EMAIL	elancon@bnl.gov
ALT EMAIL	
GITHUB ID	
INSTITUTION	Brookhaven National Laboratory

Groups and role

The screenshot shows the 'My Groups' page for the same user. The left sidebar is identical to the previous page. The main content area has tabs for MY BIO, MY INSTITUTION, and MY GROUPS (highlighted). Below the tabs is a search bar labeled 'Search My Groups'. A table lists the user's group memberships:

GROUP	ROLE	PRIVATE?	ACTION
Software and Computing		no	LEAVE GROUP
Membership Committee	Member	no	LEAVE GROUP
Elections Committee	Chair	no	LEAVE GROUP
DE&I Committee	Chair	yes	LEAVE GROUP

Below the table is a pagination control showing 'Rows Per Page' set to 10, and '1-4 of 4' items. At the bottom, there is a 'SEARCH AND JOIN GROUPS' section with a 'Search All Groups' input field.

Users can manage their group memberships
Join public groups
Leave groups

Group view – Entry to Invenio

The screenshot displays the 'ePIC COLLABORATION' web interface. The top navigation bar includes a menu icon, the site name, and user controls. The left sidebar contains various navigation links categorized under 'INSTITUTIONS', 'MEMBERS', and 'GROUPS'. The main content area is titled 'Group: Membership Committee' and features tabs for 'VIEW', 'EDIT', 'MEMBERS', 'ROLES', 'TASKS', and 'GRAPH'. The 'VIEW' tab is active, showing a form with the following details:

PARENT GROUP	Collaboration Council
NAME	Membership Committee (Collaboration Committees)
DESCRIPTION	The ePIC membership committee (eMC) is responsible for the development of the Membership policy and oversight of its fair implementation. The Membership Policy defines the process for establishing, and maintaining "Good Standing" status for Collaboration members, and the primary requirements for the authorship within ePIC. The committee will assist the Council in developing and enforcing, upon adoption of, the ePIC Membership policy.
EMAIL	epic-co-membership-committee-l@lists.bnl.gov
URL	
INVENIO	https://epic-invenio.sdcc.bnl.gov/communities/epic-membership-committee
PRIVACY	PUBLIC

Below the form is a section titled 'Members' containing a table with the following data:

ROLE	NAME	EMAIL	INSTITUTION
Vice-Chair	Pietro Antonioli	pietro.antonioli@bo.infn.it	INFN Sezione di Bologna and Dipartimento di Fisica e Astronomia, Università di
Member	Friederike Bock	bockf@ornl.gov	Oak Ridge National Laboratory
Member	Helen Caines	helen.caines@yale.edu	Yale University
Member	Robert Hancock		Test Institution

Link to document repository (InvenioRDM)

Group view – Entry to Invenio

The screenshot shows the 'ePIC COLLABORATION' web interface. The top navigation bar includes a menu icon, the text 'ePIC COLLABORATION', and a 'Finish update' button. The left sidebar contains navigation links: STATISTICS, REPRESENTATIVES, WORLD MAP, AUTHOR LISTS, MY OVERVIEW, INSTITUTIONS, FILTER INSTITUTIONS, and ADD NEW INSTITUTION. The main content area is titled 'Group: Membership Committee' and includes tabs for VIEW, EDIT, MEMBERS, ROLES, TASKS, and GRAPH. The 'VIEW' tab is active, showing details for the 'Membership Committee' group. The details include: PARENT GROUP (Collaboration Council), NAME (Membership Committee (Collaboration Committees)), DESCRIPTION (The ePIC membership committee (eMC) is responsible for the development of the Membership policy and oversight of its fair implementation. The Membership Policy defines the process for establishing, and maintaining "Good Standing" status for Collaboration members, and the primary requirements for the authorship within ePIC. The committee will assist the Council in developing and enforcing, upon adoption of, the ePIC Membership policy.), and EMAIL (epic-co-membership-committee-l@lists.bnl.gov). Below the details, there is a 'Members' section with a table listing members and their roles.

ROLE	NAME	EMAIL	INSTITUTION
Vice-Chair	Pietro Antonioli	pietro.antonioli@bo.infn.it	INFN Sezione di Bologna and Dipartimento di Fisica e Astronomia, Università di
Member	Friederike Bock	bockf@ornl.gov	Oak Ridge National Laboratory
Member	Helen Caines	helen.caines@yale.edu	Yale University
Member	Robert Hancock		Test Institution

INVENIO

https://epic-invenio.sdcc.bnl.gov/communities/epic-membership_committee

Link to document repository (InvenioRDM)

Document Management System Infrastructure

Development & Deployment

- Structured environment with production, development, and integration instances
- Automated configuration management through Puppet for consistent, reliable deployments
- Version control with Git for code development

Operations & Security

- Jira for tasks and issues tracking
- PhoneBook/COmanage-based access control
- Segregated environments separating development from production systems
- Backups for data protection and integrity
- Scalable infrastructure

Document Management System Overview

- Based on InvenioRDM v12 (the latest version)
- One-time registration, linked to ePIC PhoneBook.
- Information protection:
 - Only ePIC members can access internal records.
 - Public records are visible without authentication.
- Group and Communities:
 - ePIC PhoneBook groups = Invenio communities.
 - Group membership is managed through the PhoneBook
 - Users with a non-empty or non-'member' role in PhoneBook are considered managers of the Invenio Community (e.g., Chair, Convener, Coordinator, etc.)
 - There are two additional communities: ePIC Internal (for ePIC members only), ePIC Public (read-only for everyone)
- Any file type, up to 100MB per file, 100 files per record.

Document Management System

Some *Current* Features

- Access Control:
 - Upload, view, and comment restricted to group members.
 - Uploads require validation by group manager(s).
- Document Upload:
 - Required fields: Type, Title, Date, Contact(s).
 - Optional: Abstract, Keywords, Authors.
 - Supports drag-and-drop multiple file upload.
- Basic Workflow:
 - Managers validate uploads before publishing.
 - Email notifications are sent for submissions and comments.
- Versioning:
 - New versions can be created.
- Sharing and Public Release:
 - Records can be submitted to other ePIC groups.
 - Public visibility requires explicit public marking and manager(s) approval.

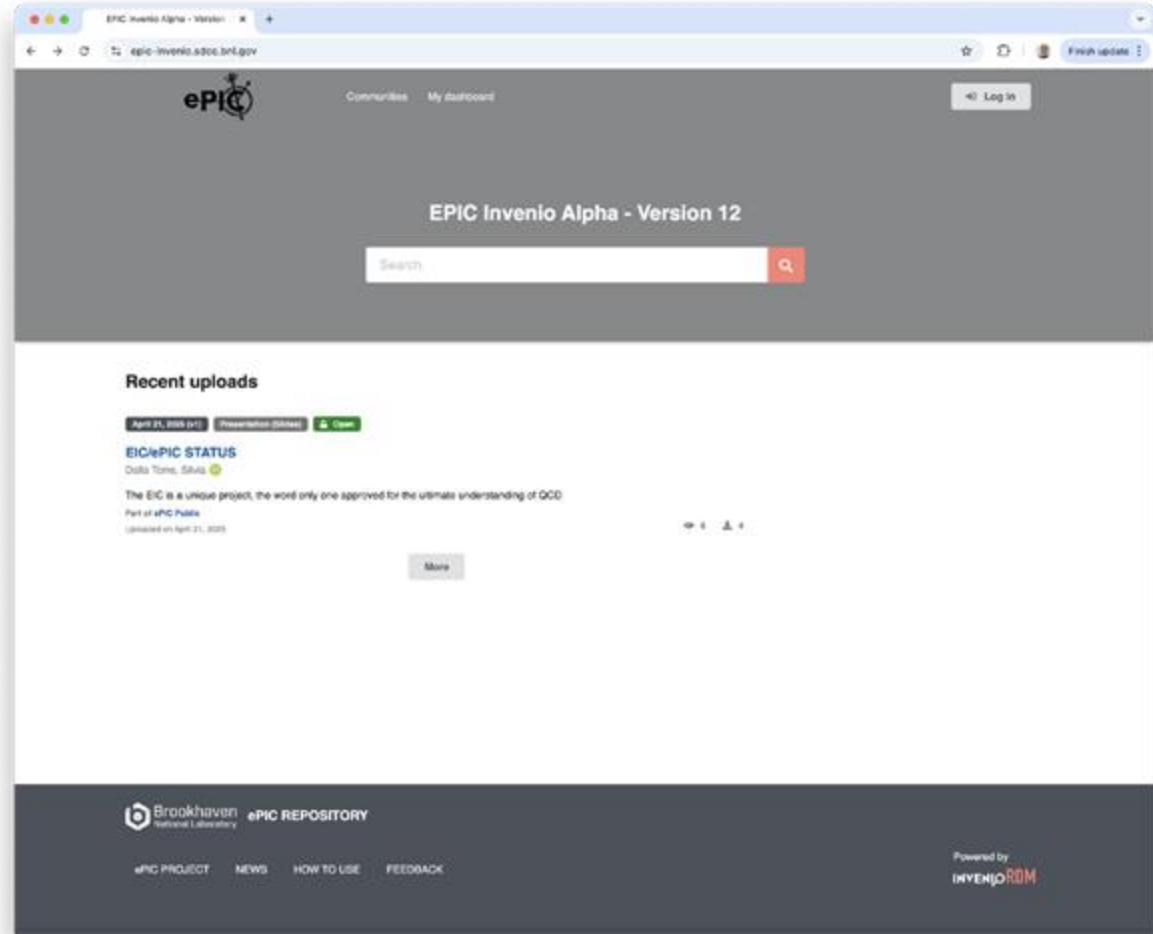
This is the current implementation

<https://epic-invenio.sdcc.bnl.gov/>

Current Layout

- InvenioRDM features a highly customizable UI.
- Branding elements (fonts, colors, logos, images) can be changed
- Menus and texts can be configured with experiment-specific fields

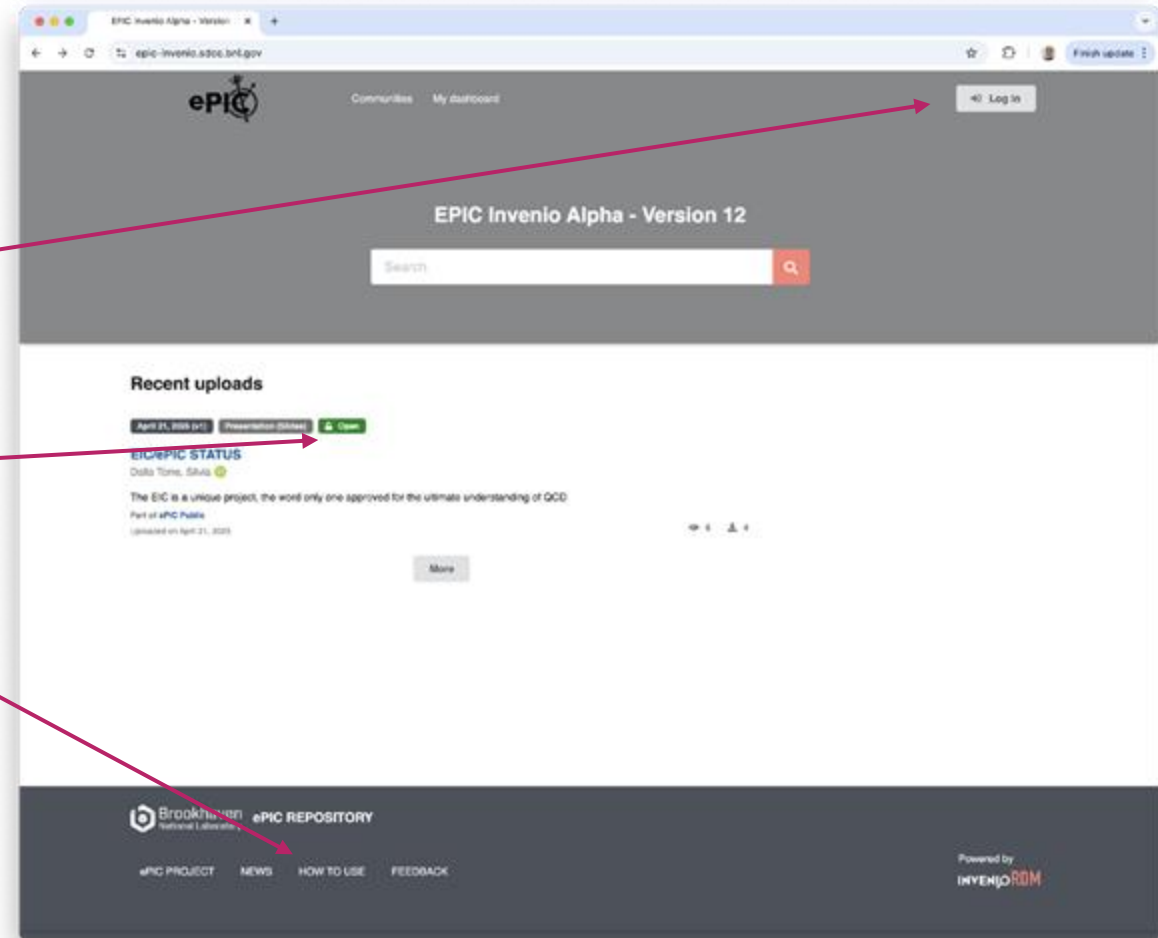
Your input is required!



<https://epic-invenio.sdcc.bnl.gov/>

Current Layout

- Authentication required to access ePIC internal documents
- Only ePIC public documents are visible on the front page
- Help (feedback) and usage instructions are accessible from the front page



This is a functional prototype, ready for testing

We actively welcome your feedback on future enhancements and customization

My groups are my communities

PhoneBook 'my groups' view

Eric Lancon
Brookhaven National Laboratory

GROUP ↑	ROLE	PRIVATE?	ACTION
Membership Committee		no	LEAVE GROUP
Elections Committee		no	LEAVE GROUP
Software and Computing		no	LEAVE GROUP
DE&I Committee	Chair	yes	LEAVE GROUP

Rows Per Page: 10 1-4 of 4

SEARCH AND JOIN GROUPS

Search All Groups

Invenio 'my communities' view

Communities

Organize, curate and collaborate on records for your institution, project, topic or event.

Search communities

+ New community

My communities [See all](#)

Internal Public Software and Computing Membership Committee Elections Committee

ePIC Internal ePIC Public ePIC: Software and Computing ePIC: Membership Committee ePIC: Elections Committee

New communities [See all](#)

There are no new communities.

Groups and communities are synchronized at login on document server

My groups are my communities

PhoneBook 'my groups' view

The screenshot shows the 'ePIC COLLABORATION' interface. The user is Eric Lancon from Brookhaven National Laboratory. The 'MY GROUPS' tab is selected. A table lists groups with columns: GROUP, ROLE, PRIVATE?, and ACTION. The first three rows are circled in red: Membership Committee, Elections Committee, and Software and Computing. The fourth row is DE&I Committee. A search bar 'Search My Groups' is at the top of the table. Below the table is a 'SEARCH AND JOIN GROUPS' section with a 'Search All Groups' bar.

GROUP	ROLE	PRIVATE?	ACTION
Membership Committee		no	LEAVE GROUP
Elections Committee		no	LEAVE GROUP
Software and Computing		no	LEAVE GROUP
DE&I Committee	Chair	yes	LEAVE GROUP

Invenio 'my communities' view

The screenshot shows the 'Communities' interface. The user is Eric Lancon. The 'My communities' section displays a grid of five communities: Internal, Public, Software and Computing, Membership Committee, and Elections Committee. The last three are circled in red. Below this is a 'New communities' section with a message 'There are no new communities.'

My communities

- Internal
- Public
- Software and Computing
- Membership Committee
- Elections Committee

New communities

There are no new communities.

Groups and communities are synchronized at login on document server

My groups are my communities

PhoneBook 'my groups' view

Eric Lancon
Brookhaven National Laboratory

GROUP	ROLE	PRIVATE?	ACTION
Membership Committee		no	LEAVE GROUP
Elections Committee		no	LEAVE GROUP
Software and Computing		no	LEAVE GROUP
DE&I Committee	Chair	yes	LEAVE GROUP

Rows Per Page: 10 1-4 of 4

SEARCH AND JOIN GROUPS

Invenio 'my communities' view

Communities

Organize, curate and collaborate on records for your institution, project, topic or event.

Search communities

New community

My communities

Internal Public Software and Computing Membership Committee Elections Committee

ePIC Internal ePIC Public ePIC: Software and Computing ePIC: Membership Committee ePIC: Elections Committee

New communities

There are no new communities.

- ePIC Internal is accessible only to ePIC members
- ePIC Public is accessible in *read-only* mode to anyone

PhoneBook groups are mapped to Invenio communities

The list of communities you belong to is updated each time you login

Group: DE&I Committee

PARENT GROUP: Collaboration Council

NAME: DE&I Committee (Collaboration Committees)

DESCRIPTION:

EMAIL:

URL:

INVENIO:

PRIVACY: PRIVATE

Members

ROLE	NAME	EMAIL	INSTITUTION
Chair	Megan Connors	meganeconnors@gmail.com	Georgia State University
Chair	Eric Lancon	elancon@bnl.gov	Brookhaven National Laboratory
Vice-Chair	Christine Natrass	christine.natrass@utk.edu	University of Tennessee, Knoxville
Member	Francesco Bossu	francesco.bossu@cea.fr	CEA-Saclay
Member	Wouter Deconinck		University of Manitoba
Member	Vincent Garonne		Test Institution

ePIC: DE&I Committee

Members

Member since	Visibility	Role
1 month ago	Hidden	Manager
1 month ago	Hidden	Reader
1 month ago	Hidden	Owner

Non-'Member' / Manager

'Member' / Reader

Communities and document handling

- Users can only upload documents to communities they belong to
- Any member of a community can view and comment on documents within that community
- All uploads require validation by community managers, defined as conveners or chairs of ePIC groups
- Records can be made of various files, for example, a note with associated figures

This is the current implementation
can be simplified

Roles and permissions

Manager and Reader roles are inherited from the PhoneBook

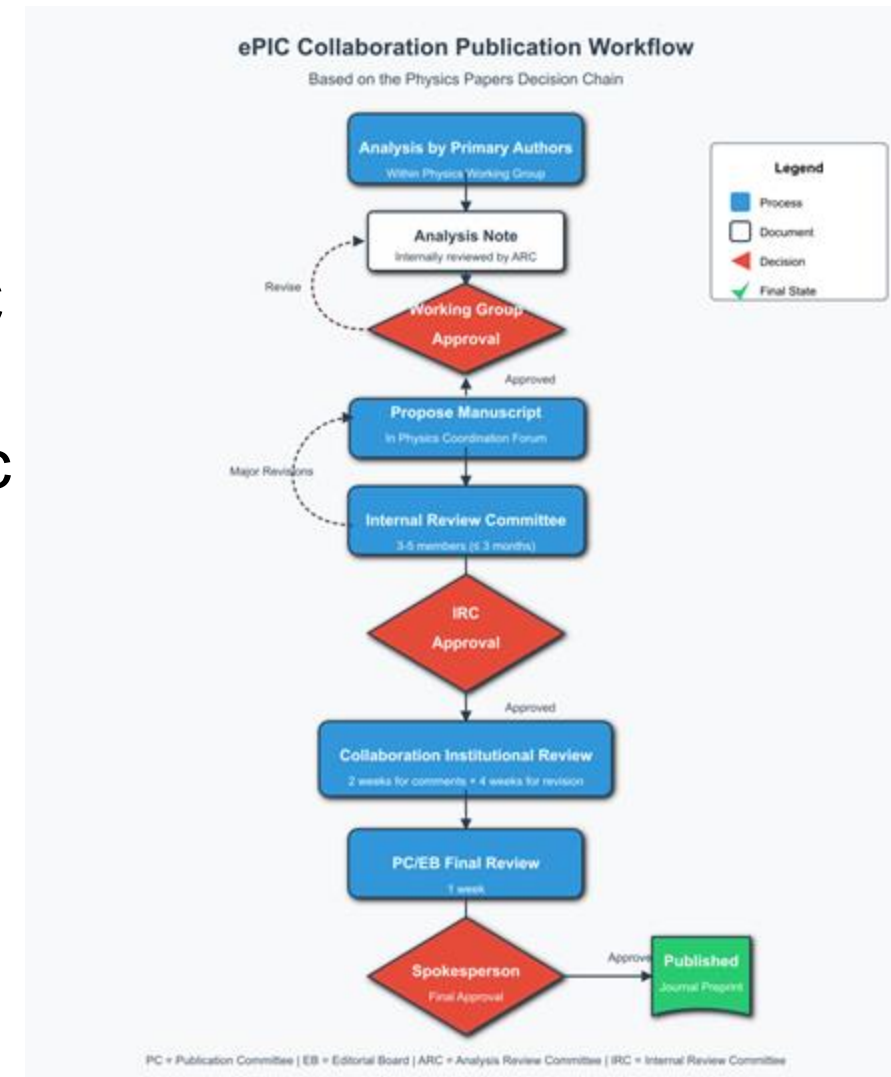
Permission	Owner	Manager	Reader
Edit community settings (title, description, logo)	✓	✗	✗
Manage roles and members	✓	✓ (limited)	✗
Accept/reject submitted records	✓	✓	✗
Publish/edit/delete community records	✓	✓	✗
View community-only records	✓	✓	✓
Comment on records	✓	✓	✓
Submit new records to the community	✓	✓	✓

ePIC should identify

- Owner(s) for ePIC communities
- Managers for 'ePIC internal' and 'ePIC public' communities

Work in progress

- Implementation of simplified workflow for efficient document management following ePIC Publication Policy
- Release of persistent identifiers (DOI) for public records
 - Integration with DataCite DOI services through OSTI infrastructure
 - Implementation of new OSTI API (released May 2024)
 - Ongoing tests
- Indico-PhoneBook integration for consistent identity management and access control



Indico PhoneBook Integration

Purpose

- Automate phonebook-to-Indico synchronization for categories and permissions
- Enable seamless global collaborative access control

Key Features

- Native Indico plugin with email-based matching
- CILogon / COnmanage authentication
- ORCID support
- Categories and user permissions based on group membership

Workflow & Status

- PhoneBook updates trigger automatic Indico updates and permission changes if needed
- Status: Pending BNL's ITD security and integration review

The screenshot displays the Indico web interface at the URL `indico-test.bnl.gov/category/523/`. The top navigation bar includes links for Home, Create event, Administration, Services, and My profile. Below this, a breadcrumb trail shows Home » Departments » ePIC. The main content area is titled "ePIC" and features a search bar labeled "Search Indico" with a "Create event" button. A sidebar on the left lists various analysis tools: Analysis Tools, Gaseous Trackers, Si Trackers, Far Backward High Rate Tracker, Forward ECAL, Barrel HCAL, Backwards ECAL, and Forward HCAL. The main panel shows the "Conference and Talks Committee" category, which is currently empty. Below this, the "Category protection" section displays a table of permissions for the category. The table lists users and their roles, with "MANAGE" and "ACCESS" buttons for each. The users listed are Brian Page, Carlos Munoz Ca..., Maria Zurek, and Nicholas Zachariou. At the bottom of the permissions table, there are tabs for "Add", "User", and "Group".

Parent Category ...	MANAGE	ACCESS
Brian Page bpage@bnl.gov	MANAGE	ACCESS
Carlos Munoz Ca... munoz@lab.org	ACCESS	
Maria Zurek zurek@anl.gov	MANAGE	ACCESS
Nicholas Zachariou nick.zachariou@york...	ACCESS	

Tailoring CRISP to ePICs Needs

- Implement feedback from ePIC collaboration members
- Establish formal contact between CRISP and ePIC
- Align development with ePIC's needs and workflows
- Document technical requirements



<https://epic-invenio.sdcc.bnl.gov/>

Login

First time
Enter a username and affiliation

Log in to account
Sign in with CILogon

Select an Identity Provider
Brookhaven National Laboratory
Remember this selection
Log On

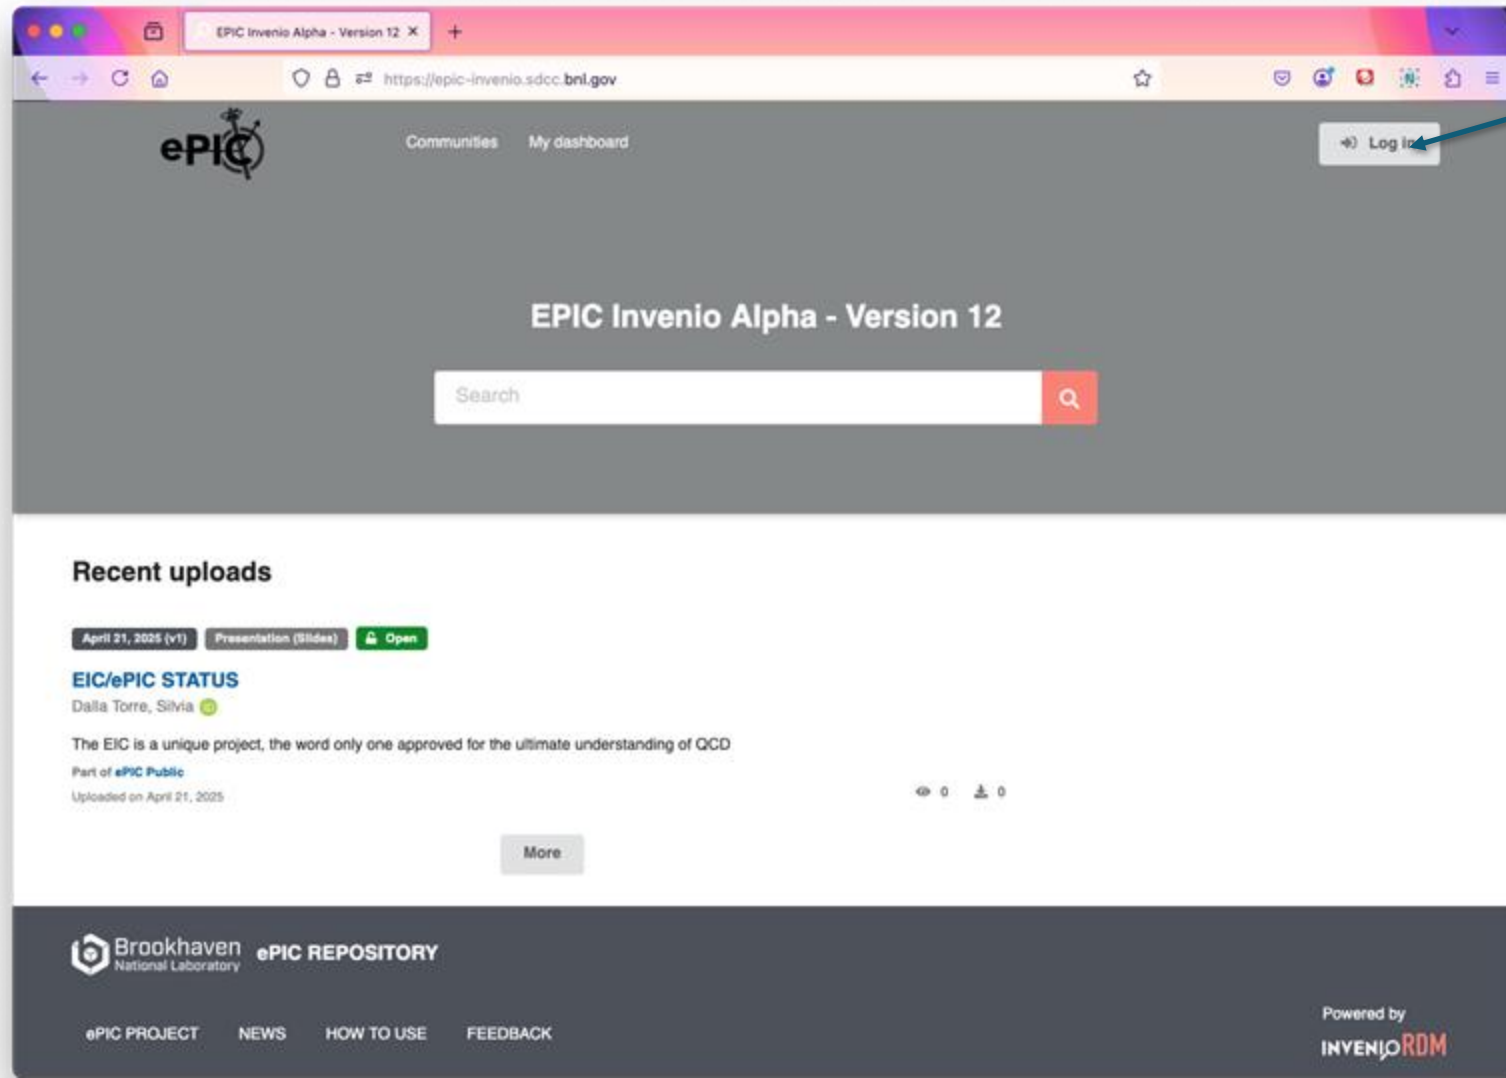
Profile
Username: elancon
Full name: Lancon, Eric
Affiliations: BNL
Email address: elancon@bnl.gov

If you are an ePIC member

You are connected

One-time registration required

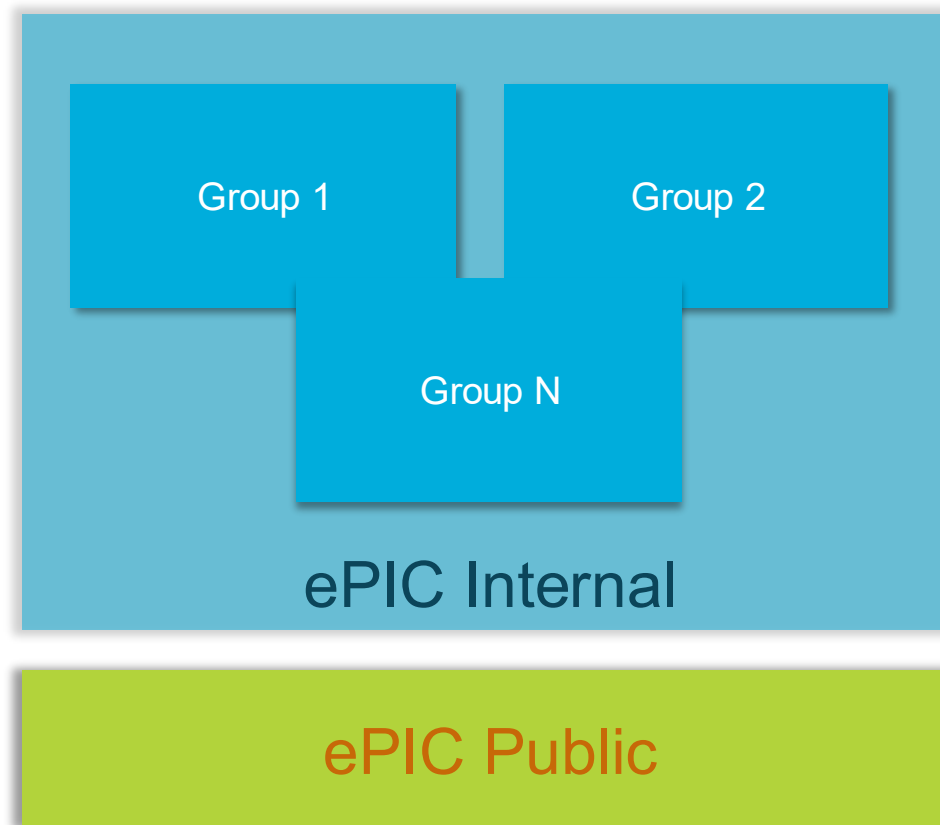
Documents marked as public on the ePIC site are visible to everyone



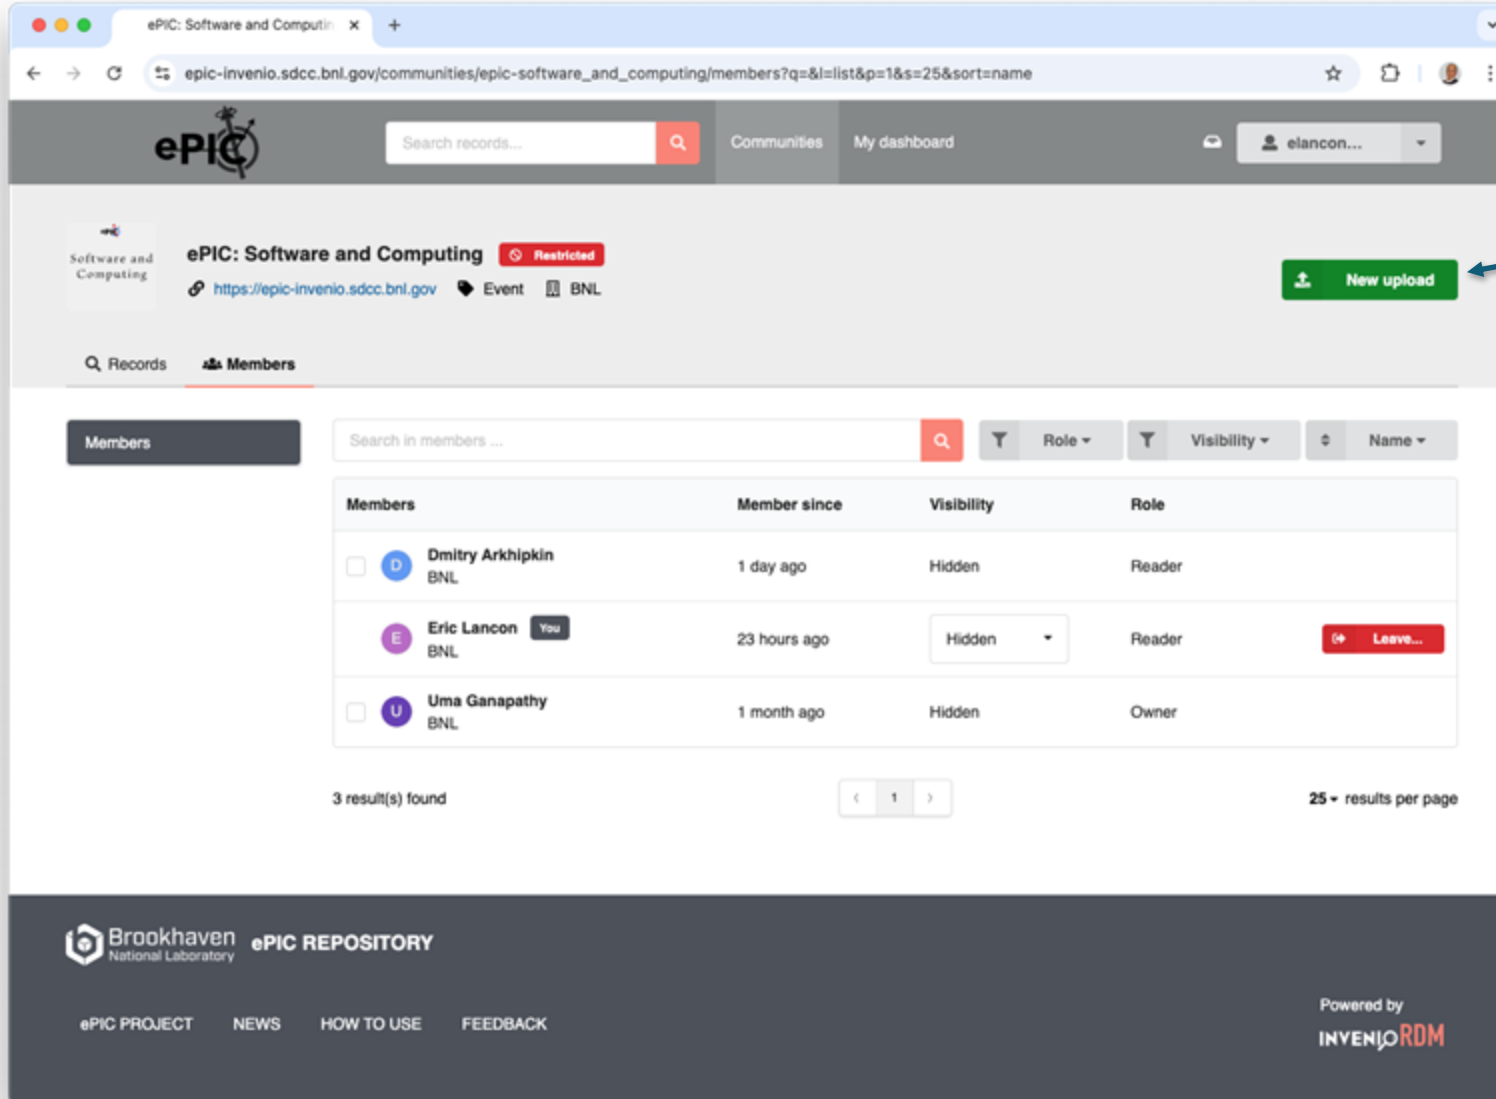
Not authenticated

ePIC groups are Invenio communities

- With two additional communities:
ePIC Internal and **ePIC Public**
 - **ePIC Internal** is accessible only to ePIC members
 - **ePIC Public** is accessible in *read-only* mode to anyone
- Community management is restricted to managers listed in the ePIC PhoneBook (group conveners, chair etc...)



Document upload to a community



The screenshot shows the ePIC Software and Computing community page. The browser address bar displays the URL: `epic-invenio.sdcc.bnl.gov/communities/epic-software_and_computing/members?q=&l=list&p=1&s=25&sort=name`. The page header includes the ePIC logo, a search bar, and navigation links for 'Communities' and 'My dashboard'. The user 'elancon...' is logged in. The main content area shows the community name 'ePIC: Software and Computing' with a 'Restricted' status and a 'New upload' button. Below this, there is a 'Members' section with a search bar and filters for Role, Visibility, and Name. A table lists three members: Dmitry Arkhipkin (Reader), Eric Lancon (Reader, marked 'You'), and Uma Ganapathy (Owner). The footer includes the Brookhaven National Laboratory logo, 'ePIC REPOSITORY', and links for 'ePIC PROJECT', 'NEWS', 'HOW TO USE', and 'FEEDBACK'. It also mentions 'Powered by INVENIO RDM'.

Search records...

Communities My dashboard

elancon...

ePIC: Software and Computing **Restricted**

<https://epic-invenio.sdcc.bnl.gov> Event BNL

Records Members

Members

Search in members ...

Role Visibility Name

Members	Member since	Visibility	Role
<input type="checkbox"/> Dmitry Arkhipkin BNL	1 day ago	Hidden	Reader
<input checked="" type="checkbox"/> Eric Lancon BNL You	23 hours ago	Hidden	Reader Leave...
<input type="checkbox"/> Uma Ganapathy BNL	1 month ago	Hidden	Owner

3 result(s) found

25 results per page

Brookhaven National Laboratory ePIC REPOSITORY

ePIC PROJECT NEWS HOW TO USE FEEDBACK

Powered by INVENIO RDM

Select New upload

Document upload - 1

The screenshot shows the 'New upload' page in the ePIC system. The browser address bar shows the URL: `epic-invenio.sdcc.bnl.gov/uploads/ftp2-g9623`. The page header includes the ePIC logo, a search bar, and navigation links for 'Communities' and 'My dashboard'. The user is logged in as 'elancon...'. The main content area is titled 'ePIC: Software and Computing' and includes a 'Change' button and a 'Remove' button. A red box highlights the file upload section, which contains a table of uploaded files, a drag-and-drop zone, and an 'Upload files' button. A red arrow points from the bullet points on the right to the 'Upload files' button.

Preview	Filename	Size	Progress
<input type="checkbox"/>	DallaTorre_EIC-ePIC_DIS-2025_CapeTown_March_2025.pdf md5:2bd9f4823e8070040941e77096520ba	14.18 MB	100%

Storage available: 1 out of 100 files, 14.18 MB out of 10.00 GB

Drag and drop files - or - Upload files

File addition, removal or modification are not allowed after you have published your upload.

Basic information

Resource type

Title

Add titles

- File upload or drag and drop
- Several files allowed

Document upload - 2

The screenshot shows the 'New upload' page in the ePIC system. The browser address bar shows the URL: `epic-invenio.sdcc.bnl.gov/uploads/ftp2-g9623`. The page header includes the ePIC logo, a search bar, and navigation links for 'Communities' and 'My dashboard'. The user is logged in as 'elancon...'. The main content area is divided into two sections. The left section, titled 'Files', contains a list of record types: Conference note, Conference proceedings, Detector paper, Image, Internal note, Performance plot(s), Physics paper, Physics plot(s), Presentation (Slides), Public note, Publication, Simulation plot(s), and Technical note. The right section, titled 'Draft', contains buttons for 'Save draft', 'Preview', 'Submit for review', and 'Share'. Below the 'Files' list, there is a 'Title' field with a red asterisk indicating it is required. A red box highlights the 'Files' list and the 'Title' field. Arrows point from the text annotations to these elements.

Files

- Conference note
- Conference proceedings
- Detector paper
- Image
- Internal note
- Performance plot(s)
- Physics paper
- Physics plot(s)
- Presentation (Slides)
- Public note
- Publication
- Simulation plot(s)
- Technical note

Conference note

Title *

Draft

Save draft Preview

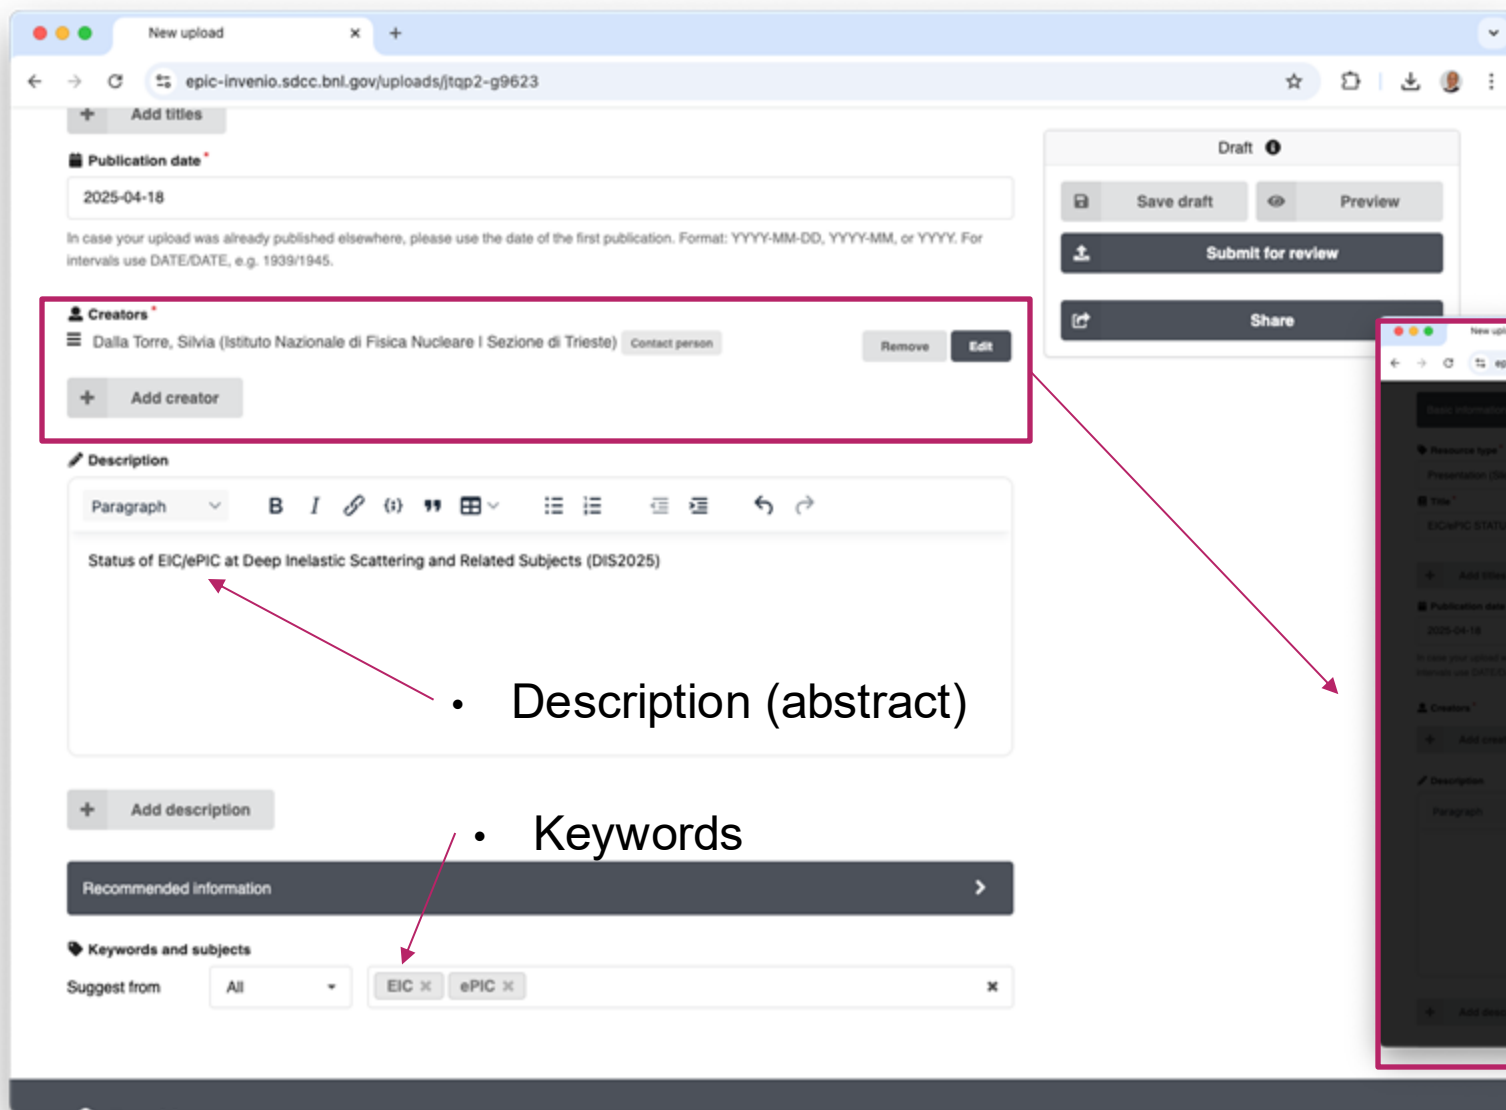
Submit for review

Share

- Select type of record
This is customizable

- Enter Title for record
** required*

Document upload - 3



epic-invenio.sdcc.bnl.gov/uploads/jtqp2-g9623

+ Add titles

Publication date *

2025-04-18

In case your upload was already published elsewhere, please use the date of the first publication. Format: YYYY-MM-DD, YYYY-MM, or YYYY. For intervals use DATE/DATE, e.g. 1939/1945.

Creators *

Dalla Torre, Silvia (Istituto Nazionale di Fisica Nucleare I Sezione di Trieste) Contact person Remove Edit

+ Add creator

Description

Paragraph

Status of EIC/ePIC at Deep Inelastic Scattering and Related Subjects (DIS2025)

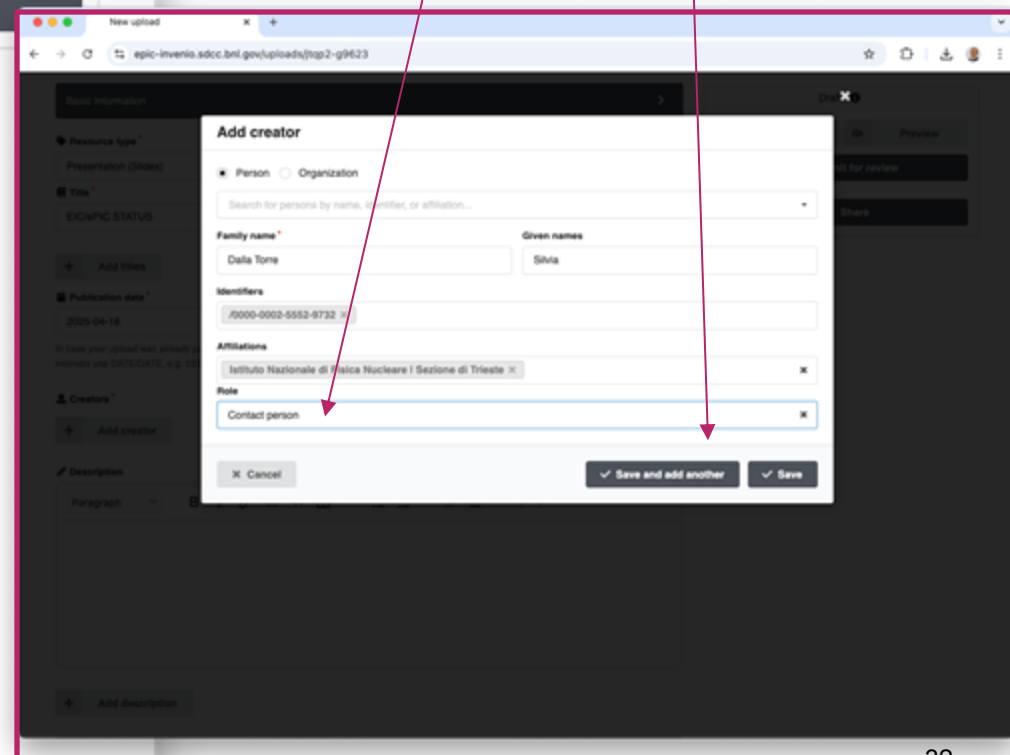
+ Add description

Recommended information

Keywords and subjects

Suggest from All EIC ePIC

- Several authors can be listed
- Each with a role
This is customizable



New upload

epic-invenio.sdcc.bnl.gov/uploads/jtqp2-g9623

Resource type *

Presentation (Default)

Title *

EPIC STATUS

+ Add title

Publication date *

2025-04-18

In case your upload was already published elsewhere, please use the date of the first publication. Format: YYYY-MM-DD, YYYY-MM, or YYYY. For intervals use DATE/DATE, e.g. 1939/1945.

Creators *

+ Add creator

Add creator

Person Organization

Search for persons by name, identifier, or affiliation...

Family name *

Dalla Torre

Given names

Silvia

Identifiers

0000-0002-5552-9732

Affiliations

Istituto Nazionale di Fisica Nucleare I Sezione di Trieste

Role

Contact person

Cancel Save and add another Save

Document upload - 4

New upload

epic-invenio.sdcc.bnl.gov/uploads/jtqp2-g9623

+ Add titles

Publication date

2025-04-18

In case your upload was already published elsewhere, please use the date of the first publication. Format: YYYY-MM-DD, YYYY-MM, or YYYY. For intervals use DATE/DATE, e.g. 1939/1945.

Creators

Dalla Torre, Silvia (Istituto Nazionale di Fisica Nucleare I Sezione di Trieste) [Contact person](#) [Remove](#) [Edit](#)

+ Add creator

Description

Paragraph **B** *I* [Link](#) [Cite](#) [Quote](#) [Table](#) [List](#) [List](#) [List](#) [List](#) [Undo](#) [Redo](#)

Status of EIC/ePIC at Deep Inelastic Scattering and Related Subjects (DIS2025)

+ Add description

Recommended information >

Keywords and subjects

Suggest from All [EIC](#) [ePIC](#) x

Draft

[Save draft](#) [Preview](#)

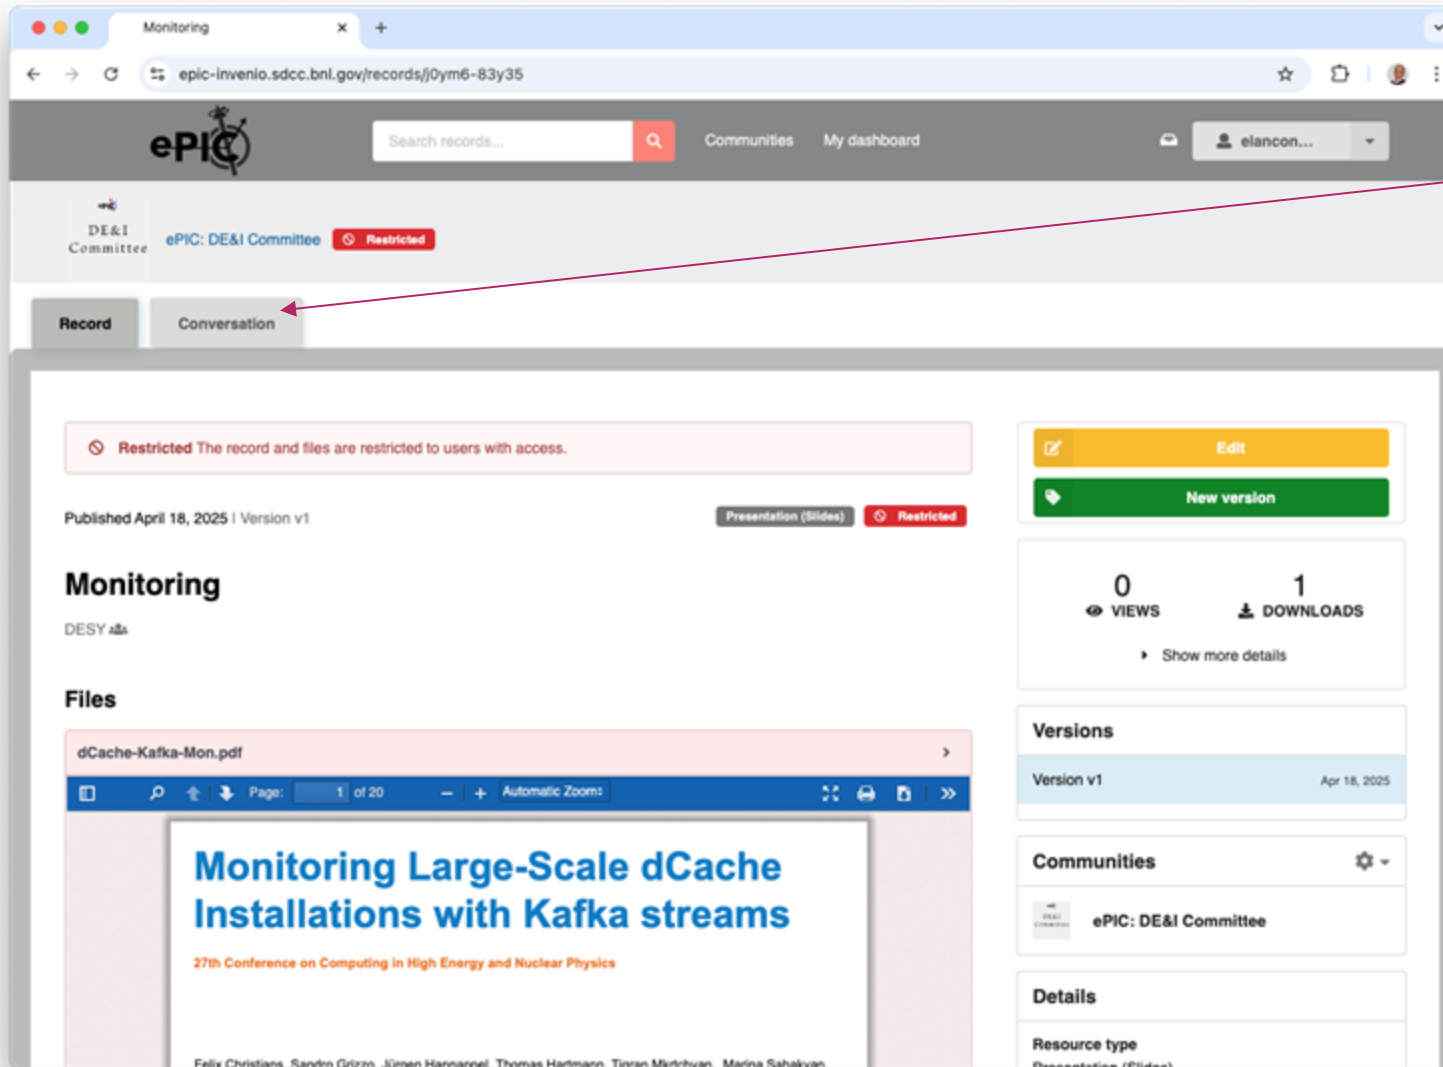
[Submit for review](#)

[Share](#)

Submit the record

Community manager(s) will receive an email and will approve or not the submission

When accepted, the record is visible within the community



- Conversations are chain of comments.
- Comments on the records can be posted by community members

Author of the submission and community manager(s) receive email on new comment

Chain of comments

Important:

For public records, always upload a fresh version of the document. Do not move a document from an internal ePIC community to the public ePIC community, as this will expose internal conversations to everyone.

Long comments can be expanded

The screenshot displays the ePIC Monitoring interface. The main window shows a conversation thread with the following comments:

- Eric Lancon** (commented 3 days ago): Hello @Garvin your request is accepted.
- Eric Lancon** (commented 3 days ago): This presentation on "Monitoring Large-Scale dCache Installations with Kafka streams" provides a ...
[Show more](#)
- Garonne, Vincent** (commented 3 days ago): Dummy comment
- Eric Lancon** (commented 3 days ago): Thank you for your answers to the comments

A red arrow points from the [Show more](#) link in the second comment to an expanded view of that comment on the right. The expanded view shows the full text of the comment:

This presentation on "Monitoring Large-Scale dCache Installations with Kafka streams" provides a comprehensive overview of how DESY uses Kafka for monitoring and managing their extensive dCache storage infrastructure. As a reviewer, here are my comments:

Strengths:

1. The presentation provides an excellent overview of dCache's significance in the high-energy physics community, establishing clear context for the monitoring solution. The statistics about dCache usage (75% of remote LHC data, usage at 9 of 13 WLCG Tier-1 centers) effectively demonstrate its importance in the scientific computing ecosystem.
2. The Kafka-based monitoring architecture is well-explained, with clear diagrams showing data flows from sources through to consumers. The step-by-step explanation of how data moves from dCache through Kafka to various analysis tools makes the system architecture accessible.
3. The practical examples of custom consumers (like the pool disabled alarming system) demonstrate real-world applications of the monitoring system. The inclusion of actual implementation screenshots adds credibility and practical value.
4. The integration with established tools (Elasticsearch, Kibana, Apache Spark) shows a thoughtful approach to leveraging existing technologies rather than reinventing solutions.
5. The presentation nicely demonstrates how Kafka enables both real-time monitoring and longer-term data analysis, providing immediate operational benefits while supporting advanced analytics.
6. The Sync&Share example effectively illustrates how the monitoring infrastructure can be extended to improve other services, showing the broader applicability of the approach.
7. The code examples for custom consumers provide actionable insights for attendees interested in implementing similar solutions.

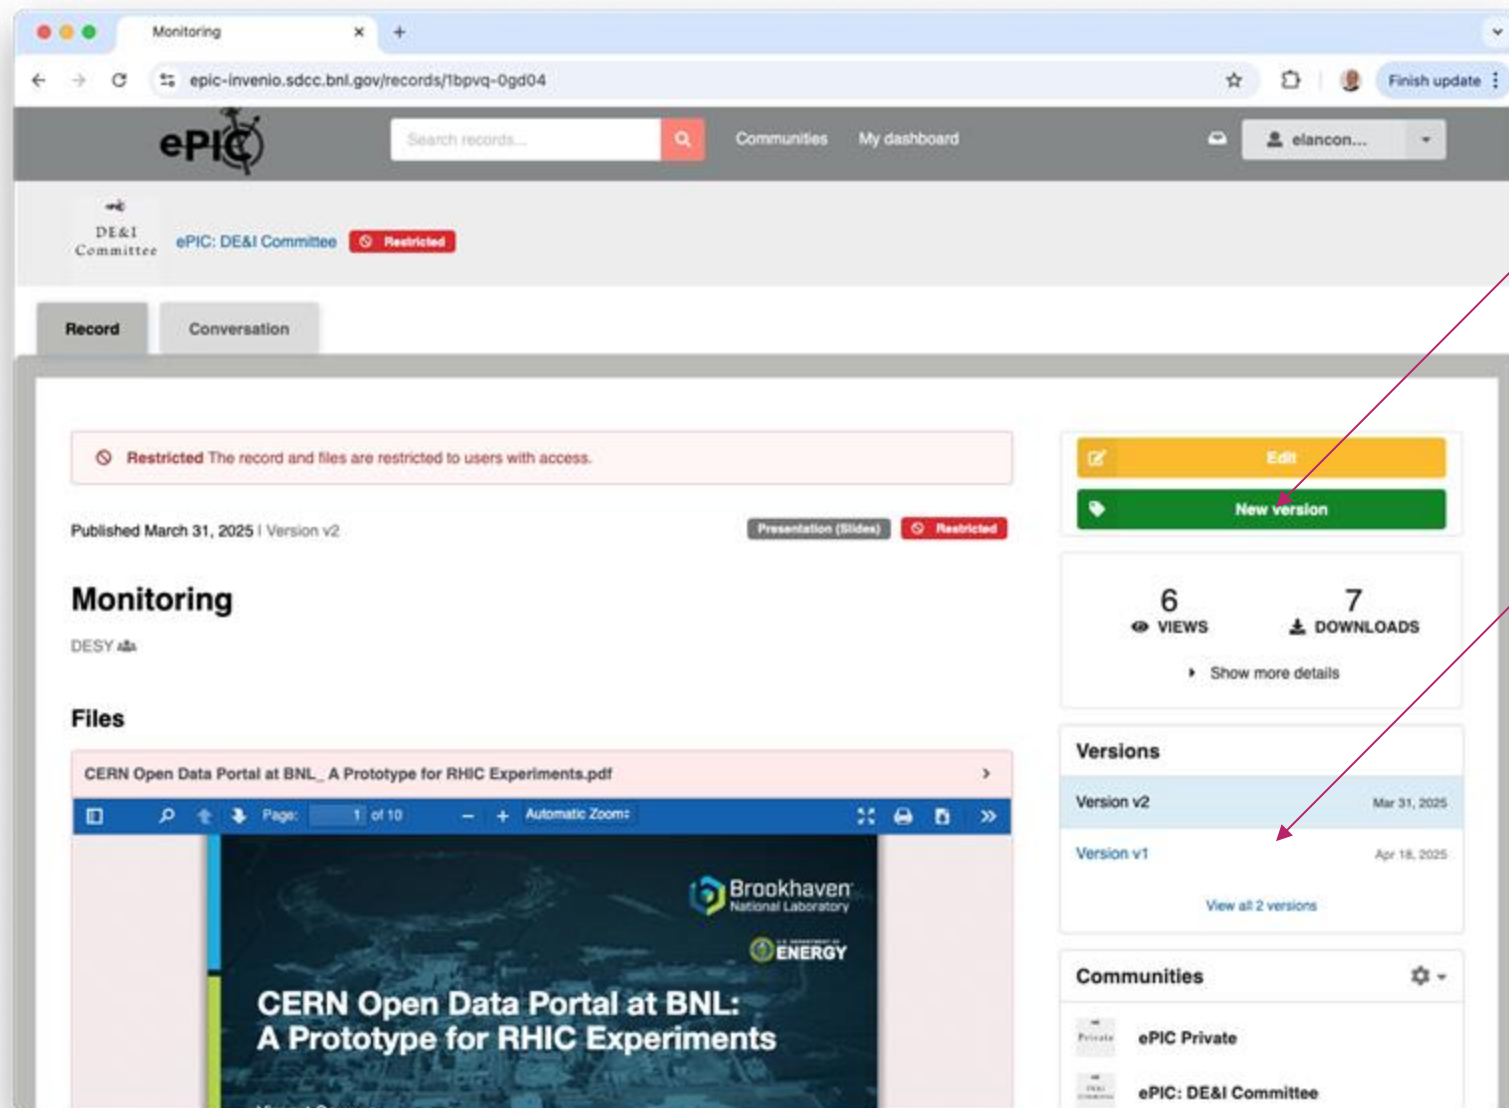
Suggestions for improvement:

1. The slides could benefit from more specific performance metrics of the Kafka implementation (throughput, latency, resource usage). While the presentation mentions producing "10-20M messages a day" with broker loads of 0.05, additional metrics would help attendees evaluate the scalability of this approach for their own environments.
2. More details on any challenges encountered during implementation and how they were overcome would add valuable perspective. Every implementation faces obstacles, and sharing these would provide realistic expectations for others considering similar solutions.
3. While the presentation mentions collaborating with BNL on ML applications, this topic could be expanded with concrete examples or preliminary results. Given the growing interest in applying ML to infrastructure management, more details here would be valuable.
4. A brief comparison with alternative monitoring approaches or technologies would help contextualize why Kafka was the right choice for DESY.

The right sidebar of the expanded view shows the following information:

- Creator:** Garonne, Vincent
- Receiver:** ePIC: DESI Committee
- Request type:** Draft review
- Status:** Accepted
- Created:** 3 hours ago
- Record:** Monitoring

Versioning of records



- New versions of a record can be created

- Access to individual versions

Records can be shared to other communities

The screenshot shows a web browser window displaying a record in the invenioRDM system. The URL is `epic-invenio.sdcc.bnl.gov/records/1bpvq-0gd04`. The record is titled "Monitoring" and is associated with the organization "DESY". It is a "Presentation (Slides)" and is marked as "Restricted". The record was published on March 31, 2025, and is version v2. The main content area shows a PDF viewer for the file "CERN Open Data Portal at BNL_ A Prototype for RHIC Experiments.pdf". The PDF cover features the Brookhaven National Laboratory and Energy logos, and the title "CERN Open Data Portal at BNL: A Prototype for RHIC Experiments" by Vincent Garonne, dated 04-17-2024. The file size is 679.6 kB. On the right sidebar, there are buttons for "Edit" and "New version". Below these, it shows 6 views and 7 downloads. The "Versions" section lists Version v2 (Mar 31, 2025) and Version v1 (Apr 18, 2025). The "Communities" section shows a list of communities: "ePIC Private" and "ePIC: DE&I". A red arrow points to the "+ Submit to community" button in the "Communities" section. The "Details" section shows the resource type as "Presentation (Slides)". The "Citation" section shows the citation as "DESY. (2025). Monitoring. invenioRDM."

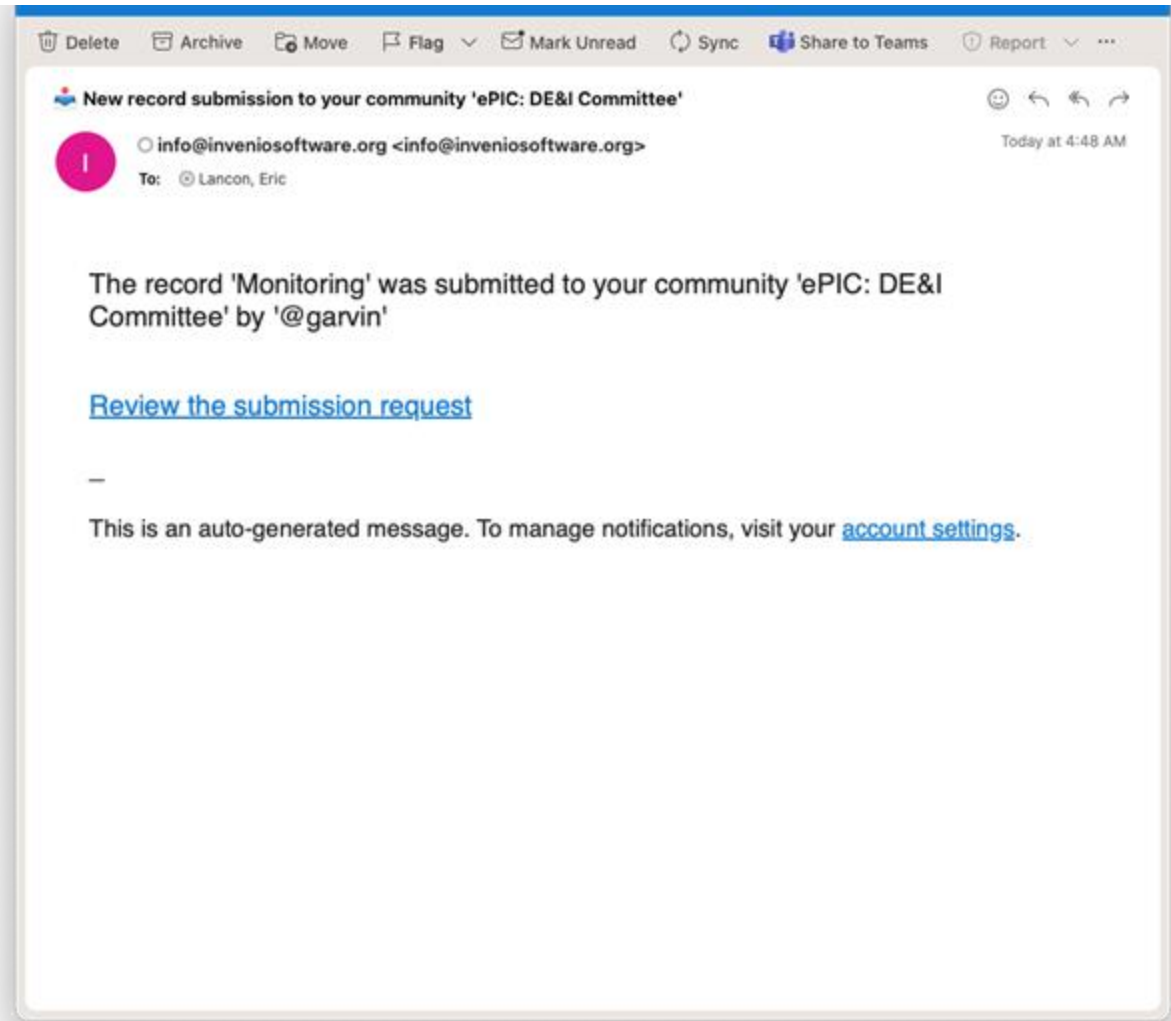
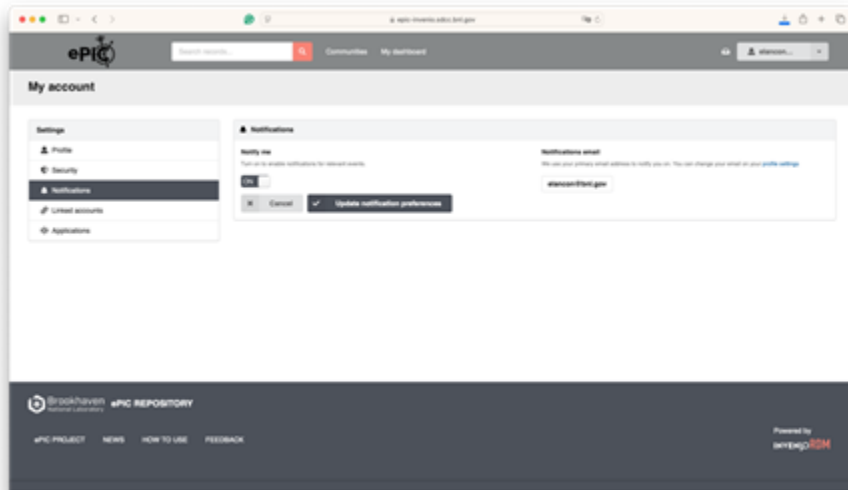
- Submit record to other community

Validation of records

User garvin submitted a record to a community

The community manager receives an email

Ensure that notifications are ON on user profile

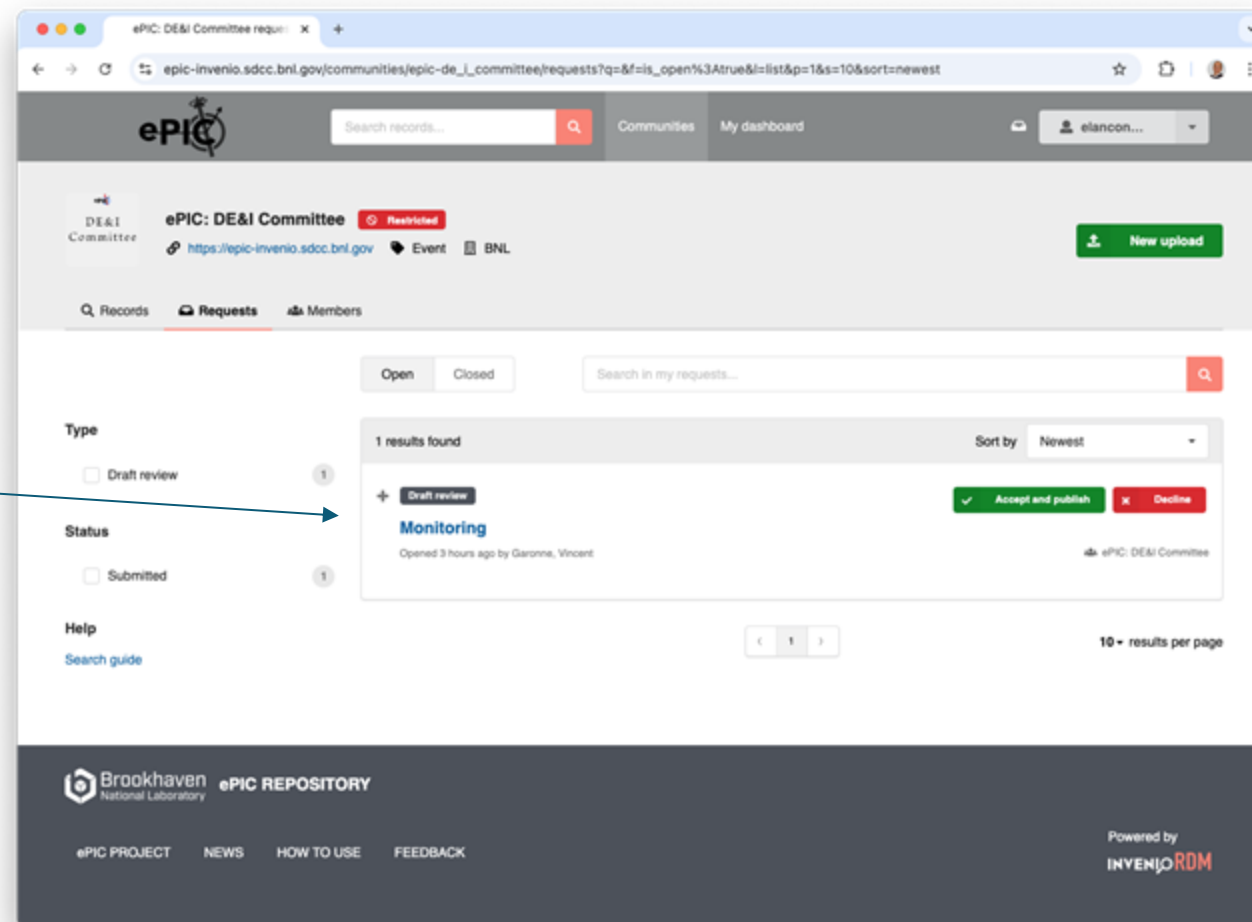


Validation of records

User garvin submitted a record to a community

The community manager receives an email

The request is visible by the manager,

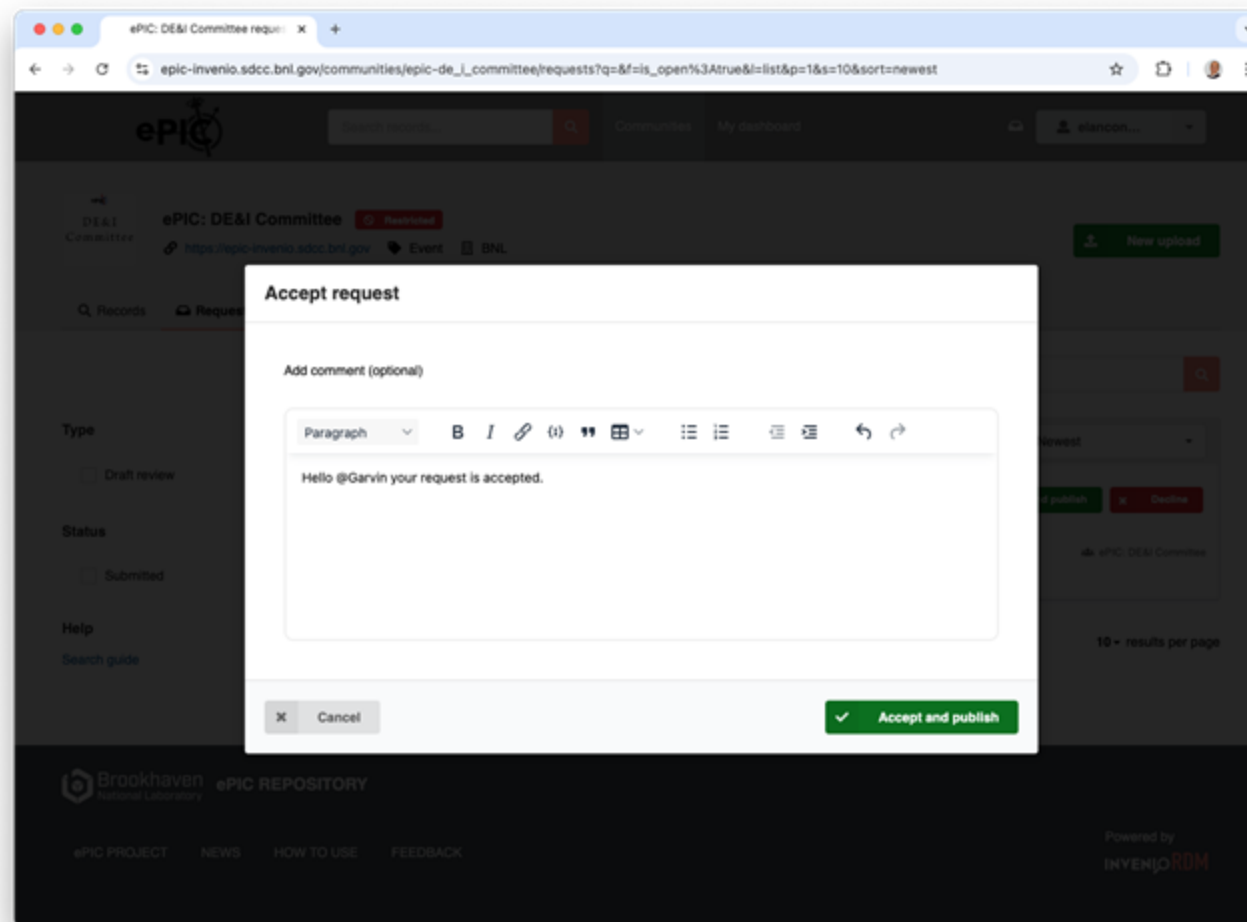


Validation of records

User garvin submitted a record to a community

The community manager receives an email

The request is visible by the manager, who can accept or decline the request

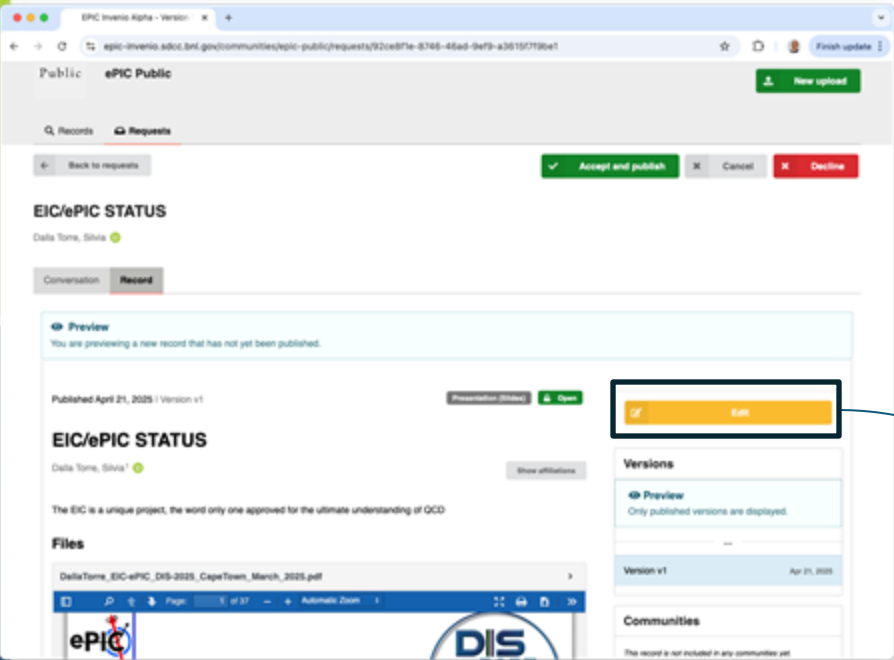


Release of public records

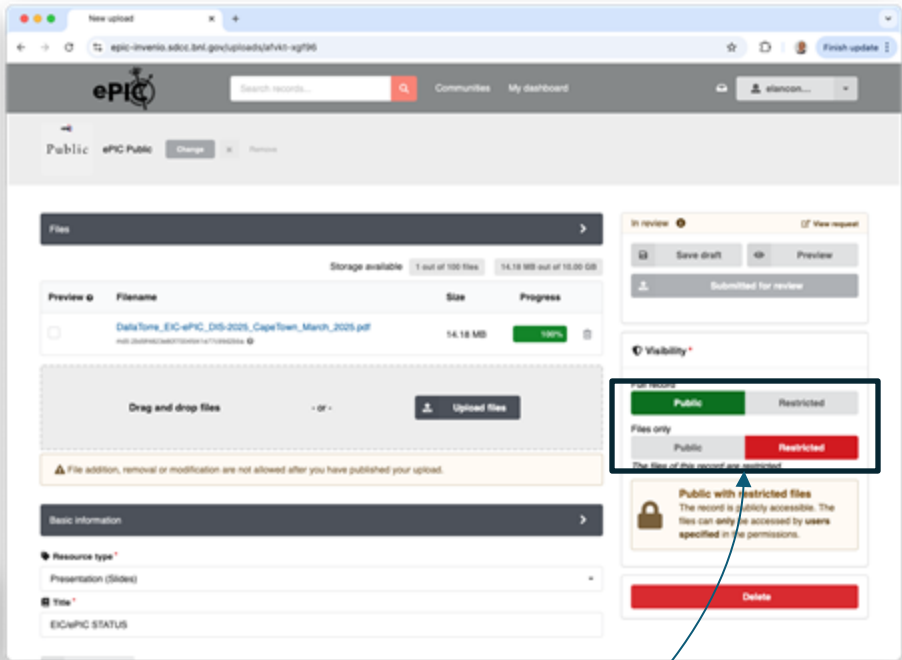
This is the current implementation

- By default, all records are restricted, and submissions must be approved by community manager(s)
- For records intended to be made public, before submission, both the record and its associated document(s) must be explicitly marked as public by the submitter.
- Without explicit public marking and manager approval, records will remain invisible to non-EPIC users, even if posted in ePIC Public.

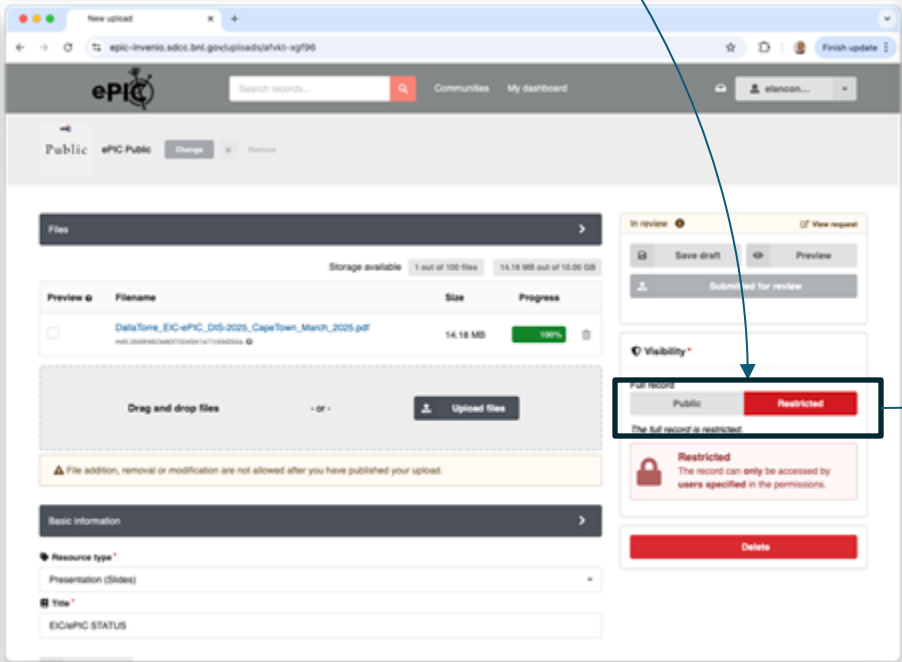
Marking documents as public



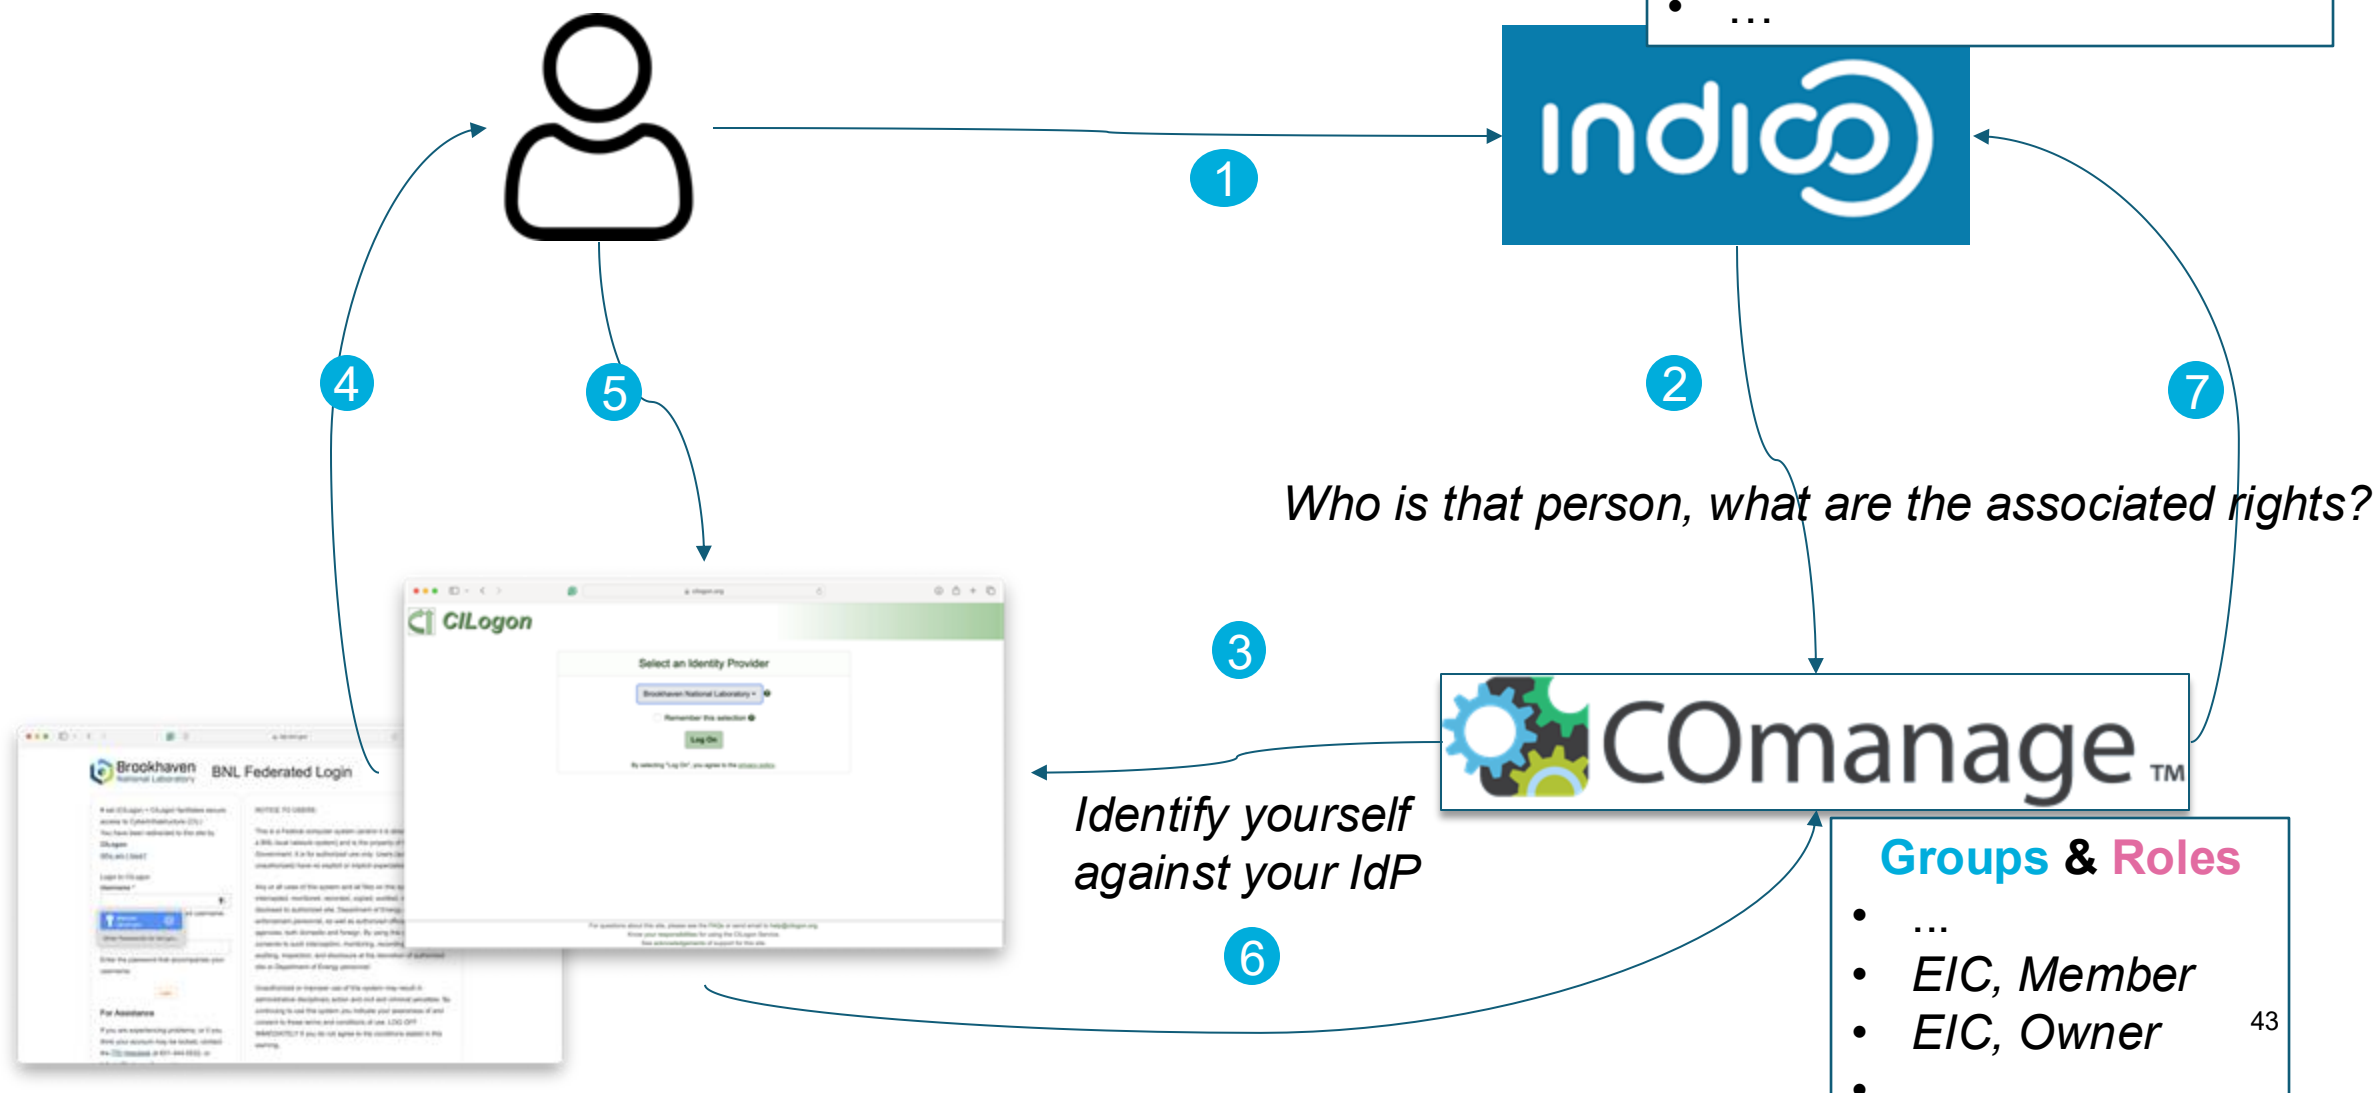
Select Edit option



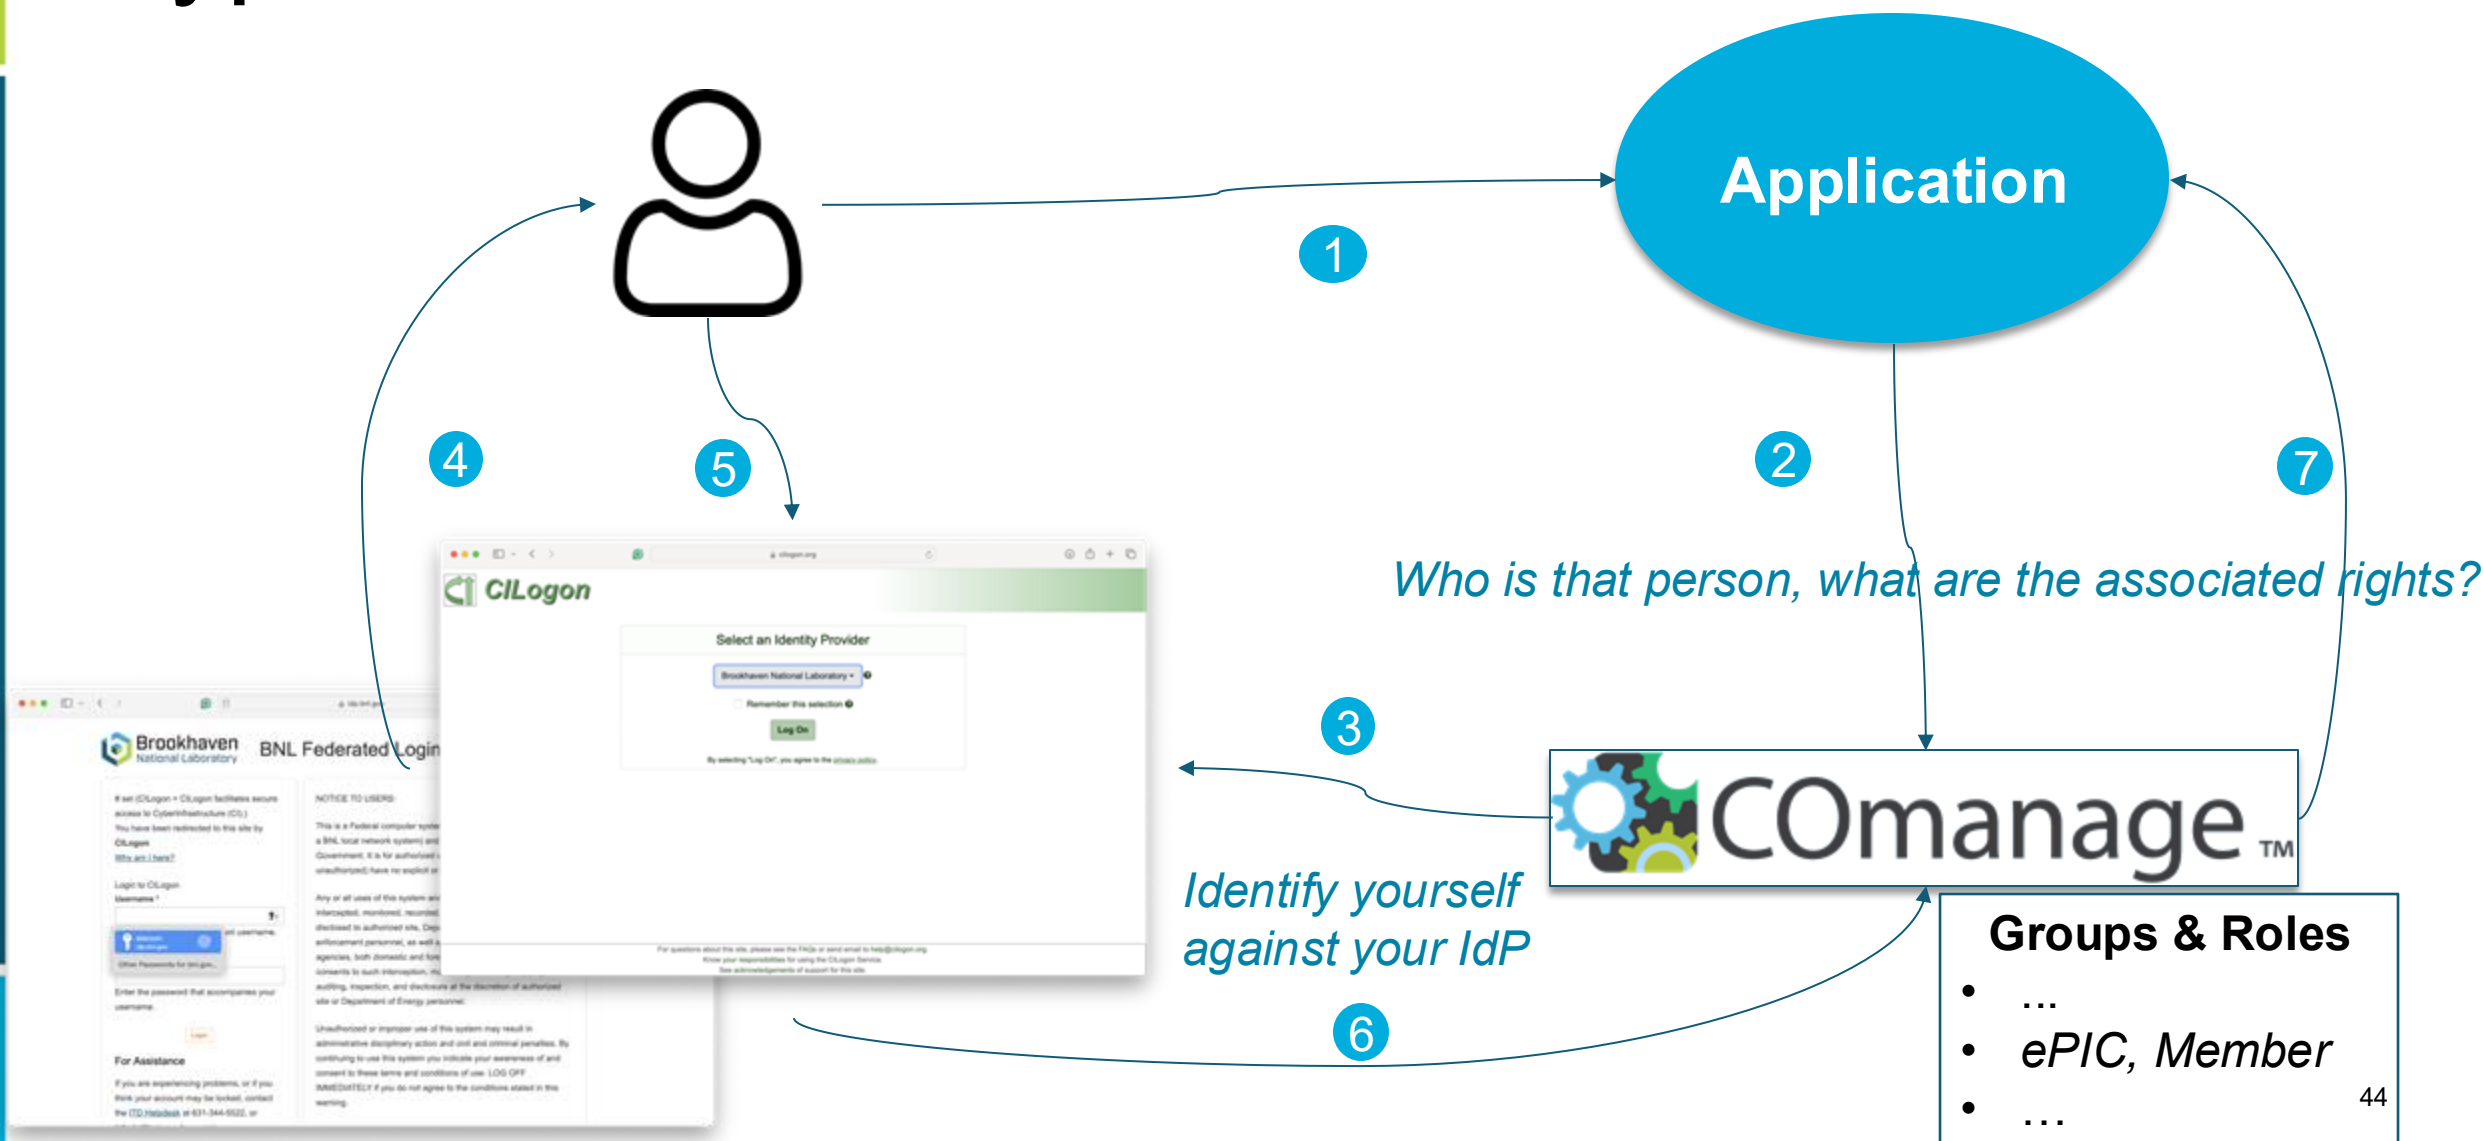
Select Public for both record and files



Mapping COnmanage Groups & Roles with Indico Categories & Permissions



Typical user workflow



PhoneBook group

Link to Invenio

The screenshot shows the 'ePIC COLLABORATION' interface. The left sidebar contains navigation links: STATISTICS, REPRESENTATIVES, WORLD MAP, AUTHOR LISTS, INSTITUTIONS (with sub-links: INSTITUTIONS, FILTER INSTITUTIONS, ADD NEW INSTITUTION, INSTITUTIONS HISTORY), MEMBERS (with sub-links: MEMBERS, FILTER MEMBERS, ADD NEW MEMBER, MEMBERS HISTORY), GROUPS (with sub-links: GROUPS, ADD NEW GROUP), and DOCUMENTS (with sub-links: DOCUMENTS, FILTER DOCUMENTS). The main content area is titled 'Group: DE&I Committee' and includes tabs for VIEW, EDIT, MEMBERS, ROLES, and GRAPH. The 'VIEW' tab is active, displaying a form with the following fields:

PARENT GROUP	Collaboration Council
NAME	DE&I Committee (Collaboration Committees)
DESCRIPTION	
EMAIL	
URL	
INVENIO	https://inv-container.sdcc.bnl.gov/communities/ePIC-de-i-committee
PRIVACY	PUBLIC

Below the form is a section titled 'Members' with a table listing the group's members:

ROLE	NAME	EMAIL	INSTITUTION
Chair	Megan Connors	meganeconnors@gmail.com	Georgia State University
Vice-Chair	Christine Natrass	christine.natrass@utk.edu	University of Tennessee, Knoxville
Member	Francesco Bossu	francesco.bossu@cea.fr	CEA-Saclay
Member	Wouter Deconinck		University of Manitoba
Member	Narbe Kalantarians	nkalandarians@vu.edu	Virginia Union University
Member	Eric Lancon - Test	elancon@bnl.gov	Brookhaven National Laboratory
Member	Iris Ponce	iris.ponce@yale.edu	Yale University

A blue arrow points from the 'Link to Invenio' text to the URL in the 'INVENIO' field of the group details form.

Group attributes

- Parent (optional)
- Name
- Description
- Sub-Category (optional)
- Email
- URL
- Status (active/inactive)
- Privacy

The screenshot displays the 'ePIC COLLABORATION' web application. The left sidebar contains navigation links for Statistics, Representatives, World Map, Author Lists, Institutions, Members, Groups, Documents, and Events. The main content area is titled 'Edit Group: Conference and Talks Committee' and includes tabs for View, Edit, Members, and Roles. The form fields are as follows:

- PARENT GROUP:** Collaboration Council
- GROUP NAME*:** Conference and Talks Committee
- GROUP DESCRIPTION*:** The Conference and Talk Committee (CTC) is responsible for the oversight and management of all oral and poster presentations given at scientific conferences on behalf of the Collaboration. This includes the selection of presenters, with
- GROUP CATEGORY:** Collaboration Committees
- GROUP EMAIL:** epic-ctc-l@lists.bnl.gov
- GROUP URL:** http://www.test.gov/
- GROUP STATUS*:** ACTIVE
- GROUP PRIVACY*:** PUBLIC GROUP

An 'UPDATE GROUP' button is located at the bottom right of the form.

Example: Coordination offices



phonebook.sdcc.bnl.gov

ePIC Collaboration

ePIC COLLABORATION

STATISTICS
REPRESENTATIVES
WORLD MAP
AUTHOR LISTS

VIEW EDIT MEMBERS ROLES

Group: Spokesperson's Office

NAME	Spokesperson's Office (Management)		
DESCRIPTION	ePIC Management		
EMAIL			
URL			
PRIVACY	PUBLIC		

Parent group

Sub-Groups

NAME	EMAIL	PRIVACY
Technical Coordination Office (Coordination Offices)		PUBLIC
Software and Computing (Coordination Offices)		PUBLIC
Analysis (Coordination Offices)		PUBLIC

Sub-groups

Members

ROLE	NAME	EMAIL	INSTITUTION
Deputy Spokesperson	Silvia Dalla Torre	silvia.dallatorre@ts.infn.it	INFN Sezione di Trieste and Dipartimento di Fisica, Universita d
Spokesperson	John Lajoie	lajoiejg@ornl.gov	Oak Ridge National Laboratory

MEMBERS
FILTER MEMBERS
ADD NEW MEMBER
MEMBERS HISTORY

GROUPS
ADD NEW GROUP

DOCUMENTS
DOCUMENTS
FILTER DOCUMENTS
ADD NEW DOCUMENT

EVENTS
EVENTS
FILTER EVENTS

Member attributes

- Firstname
- Lastname
- E-mail
- Institution
- ORCID
- Member of Council
- ...

phonebook.sdcc.bnl.gov

ePIC COLLABORATION

VIEW EDIT HISTORY

Peter Steinberg
Brookhaven National Laboratory

FIELD	VALUE	GROUP
First name	Peter	user name
Initials		user name
Last Name	Steinberg	user name
Latex Last Name		user name
Unicode name		user name
Inspire ID		user name
Gender	Male	user name
Area	Experiment	user name
Address line 1	PO Box 5000	user address
Address line 2	Brookhaven National Laboratory	user address
Address line 3		user address
City	Upton	user address
State/Region	New York	user address
Country	US	user address
Postcode/zipcode	11973	user address
Home institution	Brookhaven National Laboratory	user address
Additional Institution Affiliations	Massachusetts Institute of Technology (MIT)	user address
E-mail address	steinberg@bnl.gov	user contact information
Alternative e-mail address	steinber@gmail.com	user contact information
GitHub ID		user contact information
Work phone	631-344-4390	user contact information
Cell phone	917-443-7417	user contact information
Fax number		user contact information

STATISTICS
REPRESENTATIVES
WORLD MAP
AUTHOR LISTS
INSTITUTIONS
INSTITUTIONS
FILTER INSTITUTIONS
ADD NEW INSTITUTION
INSTITUTIONS HISTORY
MEMBERS
MEMBERS
FILTER MEMBERS
ADD NEW MEMBER
MEMBERS HISTORY
GROUPS
GROUPS
ADD NEW GROUP
DOCUMENTS
DOCUMENTS
FILTER DOCUMENTS
ADD NEW DOCUMENT
EVENTS
EVENTS
FILTER EVENTS
ADD NEW EVENT
DESCRIPTORS
INSTITUTION FIELDS
MEMBER FIELDS
INSTITUTION FIELDGR...
MEMBER FIELDGROUPS
DOCUMENT FIELDS

ePIC requirements for a document repository

Missing pieces:

- Document workflow
- Meaningful IDs (epic-note-xxx)
- DOIs

ePIC document handling: basic requirements

(Draft 1, P. Steinberg & J. Lauret)

- "Document handling" encompasses dissemination and review of ePIC documents, both internal and external
 - Internal document can include internal technical notes, internal analysis notes, conference report drafts, and publication drafts (both technical and scientific)
 - I am assuming all ePIC members can see all internal docs
 - External documents can include public notes (e.g. simulation studies not intended for publication), conference reports and publications (technical and scientific)
- Document handling requires an integrated system ("ePIC document system") which incorporates several inputs
 - A personnel database (a.k.a. phonebook) that allows bidirectional searches (associate a paper with it's authors, and an author with a set of documents)
 - Presumably documents will be important data (along with project contributions) for tracking the productivity of individuals and their groups, as part of their ePIC membership requirements
 - A document workflow database (a la ATLAS glance, which tracks document metadata, including approval status, and links to resources like document draft, analysis code, etc.)
 - InvenioRDM can already handle many of these functions
 - This system can be used to assign a meaningful ID to the document (e.g. epic-note-phys-2029-01 for a physics internal note, or epic-paper-excldiff-2030-02 for a physics paper) used to refer to it internally)
 - A document service like invenioRDM, which stores document drafts and allows for comments on individual files.
- Documents not approved for public viewing are only accessible via the document service to the ePIC collaboration
 - A layer (e.g. Comanage) will be needed to define ePIC members with access
 - Subgroup teams within ePIC should also be supported
 - No DOI is needed, since nominally internal documents are steps along the way to a public one
- Documents approved for public viewing are disseminated via the same document service
 - This service should be considered to have a lifetime starting immediately, and extending indefinitely (well beyond the conclusion of ePIC and even the EIC)
 - A DOI should be issued, which can be considered valid in perpetuity, pointing to the current location of the document even if it is migrated to a different server or even an entirely different system