

DELIVERABLE 1.9 - SYSTEM ADMINISTRATION



DECEMBER 17th 2025

biont-training.eu

contact@biont-training.eu



Co-funded by
the European Union

About this report

Grant Agreement 101100604 - DIGITAL-2022-TRAINING-02

Action Acronym BioNT

Action Title Bio Network for Training

Deliverable 1.9 System Administration

Work package 1 Training design and development

Dissemination level Public

Authors S. Di Giorgio, I. Paredes Cisneros, P. Samarakoon, H. Vela Beltran,
approved by all BioNT partners

Delivery date 2025-12-17

Consortium members

Acronym	Partner
EMBL	EUROPEAN MOLECULAR BIOLOGY LABORATORY
BIOBYTE	BIOBYTE SOLUTIONS GMBH
HPCNOW	HPC NOW CONSULTING SL
UO	UNIVERSITETET I OSLO
UB	UNIVERSITAT DE BARCELONA
ZBMED	INFORMATION CENTRE FOR LIFE SCIENCE
Rlcapacity	Rlcapacity
ALU-FR	ALBERT-LUDWIGS-UNIVERSITAET FREIBURG
EPFL	ECOLE POLYTECHNIQUE FEDERALE DE LAUSANNE

Table of Contents

About this report.....	2
Consortium members.....	2
Table of Contents.....	3
Project Overview.....	4
This workshop at a glance.....	4
Training materials, resources and advertising.....	5
Advertisement.....	6
Technical infrastructure.....	7
Workshop webpage, participant registration and selection.....	7
Infrastructure for the workshop.....	8
Zoom.....	8
Shared documents.....	9
Technical software tools, programming modules and specialised hardware.....	10
Trainees: Demographics, expertise and expectations.....	10
General information.....	10
Participation.....	16
Evaluation.....	16
Impact and outcomes.....	17
Daily feedback.....	17
Post-workshop survey.....	18
Overall feedback.....	20
Participants' self-reported improvement in skill and confidence.....	21
Conclusion.....	21

Project Overview

The BioNT consortium is dedicated to providing a comprehensive training program and fostering a community for digital skills relevant to the biotechnology industry and biomedical sector. With a curriculum tailored for both beginners and advanced professionals, BioNT aims to equip individuals with the necessary expertise in handling, processing, and visualising biological data, as well as utilising computational biology tools. Leveraging the consortium's strong background in digital literacy training and extensive network of collaborations, BioNT is poised to professionalise life sciences data management, processing, and analysis skills.

This workshop at a glance

The final BioNT training “System Administration for HPC” was structured as a two-stage initiative, consisting of a pre-workshop industry meetup followed by the main three-day training workshop.

Pre-workshop industry meetup

The industry meetup event, titled “[Sysadmins in HPC meet-up](#)”, held on the 10th of October at the Barcelona Supercomputing Center (BSC), was designed to engage our target audience and raise awareness of the BioNT project, in particular, the upcoming training course “System Administration for HPC.” The half-day event brought together systems administrators, HPC practitioners, startup teams, trainers, and systems engineers interested in building or strengthening their skills in High Performance Computing operations. Through technical talks by HPC experts, an interactive panel with live Q&A, an afternoon networking session, and a guided tour of the MareNostrum 5 supercomputer, the meetup provided a practical forum to exchange knowledge, share operational experiences, and discuss current challenges and opportunities in HPC system administration. The event also served as a launchpad to present “System Administration for HPC” as a targeted BioNT training opportunity for both new and experienced sysadmins, helping them structure their learning path in this field. Participants were encouraged to provide feedback on training needs and priorities, which will be used to further adapt and refine the contents of the course.

Key discussion topics of the industry meetup event:

- Operational best practices for managing modern HPC systems
- Automation, configuration management, and monitoring for scalable HPC operations
- Skills, career paths, and onboarding strategies for new HPC system administrators
- Collaboration between research centres, industry, and startups to advance HPC infrastructure and services
- BioNT’s training offerings, with a focus on “System Administration for HPC” as a competence-building opportunity

Main training workshop

The subsequent three-day workshop was designed to equip system administrators and technical staff with the practical skills needed to deploy, manage, and maintain high-performance computing (HPC) clusters efficiently. Covering essential topics from foundational Linux administration to advanced orchestration and application building, participants gained a comprehensive overview of the core components that power modern HPC environments.

The “System Administration for HPC” workshop was structured as a three-day, hands-on training designed to build core competencies in operating and managing HPC systems. The first day focused on establishing a solid foundation in Linux system administration, covering user and group management, file permissions, filesystems (including partitions and NFS), package managers (such as `dnf` and `apt`), firewalls, and service management with SystemD. The day concluded with an introduction to cluster management concepts using Warewulf, providing a smooth transition from single-node administration to managing full HPC clusters.

The second day was dedicated to the Slurm workload manager and container technologies that underpin modern scientific workflows. From an administrative perspective, participants explored Slurm’s architecture, user and account management. This was complemented by practical sessions on Docker, Docker Compose, and Singularity, enabling attendees to understand how containers can be integrated into HPC environments to support reproducible and portable workloads.

The final day shifted the focus towards automation, monitoring, and scientific software environments. Participants used Ansible to orchestrate deployments and configuration management, and learned how to expose and visualise cluster metrics using Prometheus and Grafana. The workshop concluded with sessions on building and managing scientific application stacks using EasyBuild integrated with EESSI. By the end of the workshop, participants had gained practical experience and actionable knowledge to confidently manage HPC cluster operations and support diverse scientific software needs within their organisations.

Feedback from participants indicated exceptionally high satisfaction levels across all measured dimensions (as detailed in the section “Post-workshop Survey”). The workshop’s overall success was underscored by recommendation likelihood scores: 95% participants rated it between 8 and 10, with 58% participants giving the maximum score of 10, indicating they would very likely recommend the workshop to colleagues.

As further detailed in later sections of this report, the workshop successfully delivered high-quality instruction in a supportive learning environment with strong practical relevance for participants.


Training materials, resources and advertising

As this “System Administration for HPC” workshop was delivered for the first time, all lecture materials and hands-on exercises were developed from scratch. The content was tailored to address real-world HPC system administration scenarios, focusing on topics such as Linux fundamentals, Slurm, containers, automation, monitoring, and scientific software environments. In line with BioNT’s mission of bridging academic and industrial practices, the materials were designed to be directly applicable to production-like HPC environments. All resources are openly available on Zenodo and licensed under a Creative Commons license, enabling their reuse, adaptation, and integration into future training activities by other communities.

Date	Topic
2025-10-21	<ul style="list-style-type: none">• Introduction• Linux Administration• Introduction to Cluster Managers: Warewulf
2025-10-22	<ul style="list-style-type: none">• Slurm Administration• Container technologies: Docker
2025-10-23	<ul style="list-style-type: none">• Orchestration - Ansible• Metrics and Monitoring: Prometheus and Grafana• Scientific Application Building: EasyBuild and EESSI

Advertisement

The meet-up and the workshop were advertised via social media, several websites, mailing lists and Slack spaces of networks and communities (ELIXIR, NFDI4Microbiota and de.NBI, among others). For the advertising, a tailored image was generated, which included a QR code to facilitate access to the registration platform, as shown in Figure 1.



BioNT
958 followers
3mo · Edited ·

REMINDER: Calling future HPC experts! 📢
Join us and our experts from **Do IT Now** at our next **#BioNT** workshop: System Administration for HPC

📅 Save the Dates: October 21-23, 2025
🕒 Workshop Schedule: 09:00 - 16:00 (CEST)
📅 Registration Deadline: October 14, 2025

This 3-day workshop offers essential, practical skills to confidently administer and support HPC clusters:

- Day 1: Core Linux system administration – users, filesystems, services, firewalls
- Day 2: HPC essentials – Slurm, Docker, Singularity
- Day 3: Automation & monitoring – Ansible, Prometheus, Grafana, EasyBuild, EESSI, Spack

👤 Meet the organizing team: **Silvia Di Giorgio, Helena Vela, Teresa Müller and Susana Hernández**; get to know our helpers: **Isabela Paredes Cisneros, Pubudu Samarakoon**; and connect with the rest of the BioNT team: **Lisanna Paladini, Renato Alves, Alan O'Caïs, Ivica Letunic, Konrad Förstner, Rabea Müller, Tili Sauerwein, Tanja Ninkovic, Sabry Razick**, and **École polytechnique fédérale de Lausanne, EPFL**.

We look forward to receiving your applications for this training workshop, which is supported by **ELIXIR**, and **de.NBI**.

Visit our event webpage 🌐 to discover the program details, meet the organising team, and submit your application. Remember, participation in the workshop is completely free

#LearningOpportunity #BioinformaticsWorkshop #DigitalLifeSciences
#NetworkForTraining #ComputationalLifeSciences

🌐 <https://lnkd.in/gZBdWQIN>

BioNT ONLINE WORKSHOP
SYSTEM ADMINISTRATION FOR HIGH PERFORMANCE COMPUTING

Location: Online
Dates: October 21-23, 2025
Free of charge (registration required)

We're excited to see you there!



BioNT
958 followers
2mo ·

Sysadmins in HPC Meet-Up 📢
Are you a systems administrator in HPC or curious about stepping into this exciting world? Join us and **Do IT Now** for a half-day meet-up at the Barcelona Supercomputing Center!

📅 Date: 10 October 2025
📍 Location: Barcelona Supercomputing Center
🆓 Free of charge (registration required)

👉 What to expect:

- Technical talks from HPC professionals
- Live interactive panel with Q&A
- Networking opportunities with peers and SMEs
- Guided tour of MareNostrum 5

Whether you're a sysadmin, part of a startup, a professional seeking new challenges, or a student eager to learn, every expertise level is welcome in this event. See you in Barcelona!!!

📍 Register now 🌐
<https://lnkd.in/gkhmXzB9>

BioNT MEET-UP
SYSADMINS IN HPC MEET-UP

Location: Barcelona Supercomputing Center
Dates: October 10, 2025
Free of charge (registration required)

Link for registration

Figure 1 - Meet-up (left) and Workshop (right) of the workshop advertised through LinkedIn.

Technical infrastructure

The workshop was scheduled over 3 days with sessions delivered from 09:00 to 16:00 CEST from October 21st to October 23rd, 2025, and required registration through the event website. The entire event was conducted virtually and offered at no cost to participants.

Workshop webpage, participant registration and selection

The CECAM event management platform, provided by the EPFL, was used to create a dedicated webpage for the workshop. This webpage included the workshop description, learning objectives, requirements, program, and other relevant information for potential participants. This event webpage with the necessary information is accessible at: <https://www.cecarn.org/workshop-details/system-administration-for-hpc-1465>, where the full program, learning outcomes and requirements are listed.

Registration was required, and the CECAM platform was used to manage the applicant's information and communication. In parallel, the EMBL servers were used to collect pre- and post-workshop information through pseudo-anonymised surveys. Survey responses were linked to applicant data only via a unique identifier, which was provided during the CECAM registration process and entered again in the EMBL-based survey. This approach ensured that only the workshop organisers had access to identifiable applicant data while still collecting information relevant to the workshop separately. To register, applicants had to: (i) register on the CECAM platform, (ii) complete and submit the pre-workshop survey, and finally (iii) complete the application on the CECAM platform using the unique identifier provided in the pre-workshop survey.

The workshop, which focused on hands-on practices, required participants to use virtual machines equipped with virtualisation capabilities. Due to high costs, course organisers could only provide 30 VMs, necessitating a selection process from the applicants. The selection process was based on the answers provided in the pre-workshop survey (containing no personal information). For the selection criteria, learners currently working in European countries who were either job seekers, small-and medium-sized enterprise (SME) employees, SME collaborators, or individuals aiming to work in SMEs were prioritised. Among applicants meeting these criteria, from 45 selected participants, 30 were finally allocated on a first-come-first-served basis to ensure fair and transparent allocation of the limited computational resources.

The 30 final participants represented mostly European countries, as given by the selection criteria explained above. Career-wise, the cohort was dominated by experienced and early-career academics, with postdoctoral researchers and graduate students forming the majority, alongside research staff and analysts. Regarding SME engagement, an overwhelming number of applicants expressed aspirations to work in SMEs (SME job seekers), while the rest either collaborated with SMEs or were current SME employees. Selected participants showed strong alignment with system administration and operational knowledge on HPC clusters. This composition highlighted the selection process's effective targeting of the desired demographic, aligning with the project's goals.

Infrastructure for the workshop

Zoom

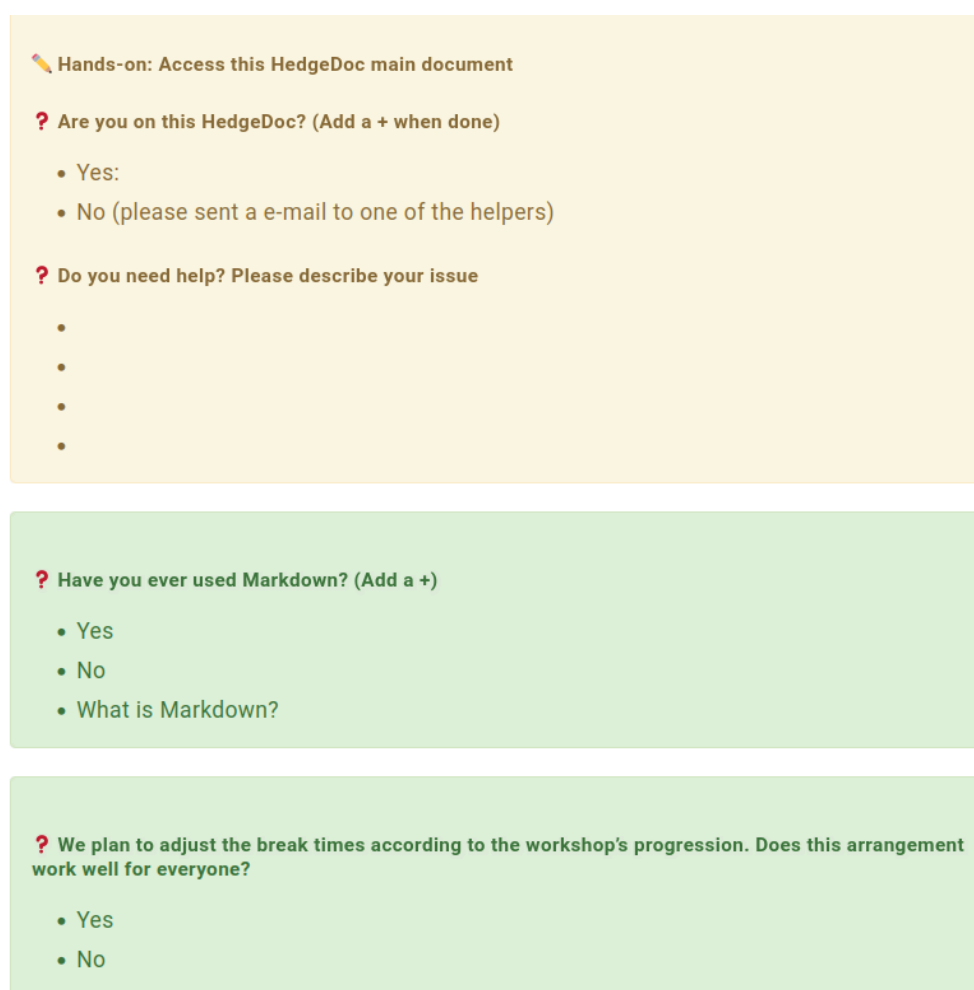
The workshop was delivered via Zoom webinar, allowing participants to learn directly from the trainers. To facilitate video recording for future self-paced learning materials, participants' microphones were muted, and cameras were off during presentations. All direct interactions were instead conducted in written form via a prestructured collaborative document, which allowed for anonymous participation. To serve the collaborative documents, a [HedgeDoc](#) collaborative space was set up by BIOBYTE, and was hosted on their server.


To further enhance the interactive experience, the instructors occasionally engaged in dialogue around the topics under discussion. The idea was to bring in some of the “real


world” experience of the instructors when it came to the topics, and to show that the simplified examples of the hands-on can have a practical impact in larger projects.

Shared documents


A [Main](#) collaborative document, set up by the instructors and organisers, was shared with the participants before the workshop. Each section of the workshop had dedicated *Hands-on* boxes to report on the task status, ask questions or raise issues. Helpers engaged and assisted participants by answering the questions and issues directly in this document. This Main document was updated live during the workshop. Separate boxes to answer questions were used to improve participant engagement and as an indirect learning assessment (Figure 2).




 Hands-on: Access this HedgeDoc main document

 Are you on this HedgeDoc? (Add a + when done)

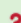
- Yes:
- No (please sent a e-mail to one of the helpers)

 Do you need help? Please describe your issue

-
-
-
-

 Have you ever used Markdown? (Add a +)

- Yes
- No
- What is Markdown?

 We plan to adjust the break times according to the workshop's progression. Does this arrangement work well for everyone?

- Yes
- No

Figure 2 - Example of a question and hands-on boxes in the collaborative document.

To help with the organisation, four HedgeDoc documents were used: (i) a [Template](#) with all instructions and boxes for hands-on, questions, etc; (ii) the [Main](#) document with the information for the participants during the workshop, filled with boxes related to the section covered by the instructor to help with the navigation and cleaned during each break to avoid an overcrowded document; and (iii) a [History](#) document collecting all the content from the Main document. This document was shared with participants during the workshop to grant them access to all prior conversations. In addition, (iv) a document for Helpers & Instructors was created, providing the workshop setup, interactions and explaining tasks of the helpers and instructors.

There was a very lively interaction with the learners through the shared document. This was made possible by the extended set of helpers available to answer participant questions.

Instructors, from time to time, used the shared document to check the satisfaction with lesson delivery and made necessary adjustments. Type-along sessions, where learners tried to follow the instructor on their computers, and demo sessions, where learners watched the instructor complete tasks, were clearly communicated to the learners.

Technical software tools, programming modules and specialised hardware

The workshop required at the beginning only the VMs to be available to the users, as they would build up their cluster through the hands-on sessions. The credentials of those were provided through an [attendance confirmation](#), where the IP and the key for accessing each VM were specified. Those were shipped with a Linux environment and the virtualisation set up for each trainee.

Trainees: Demographics, expertise and expectations

Following the selection process, we extracted the pre-workshop survey responses provided by the selected 45 participants. These responses were then analysed and shown in the following sections. The pre-workshop survey comprised 22 questions covering skills, demographics, and miscellaneous topics. The majority of questions were optional for the successful completion and submission of the survey.

General information

Of the 45 first selected participants, 30 were male, 13 female, and 2 preferred not to specify. Nationality and employment are summarised in Figure 3.

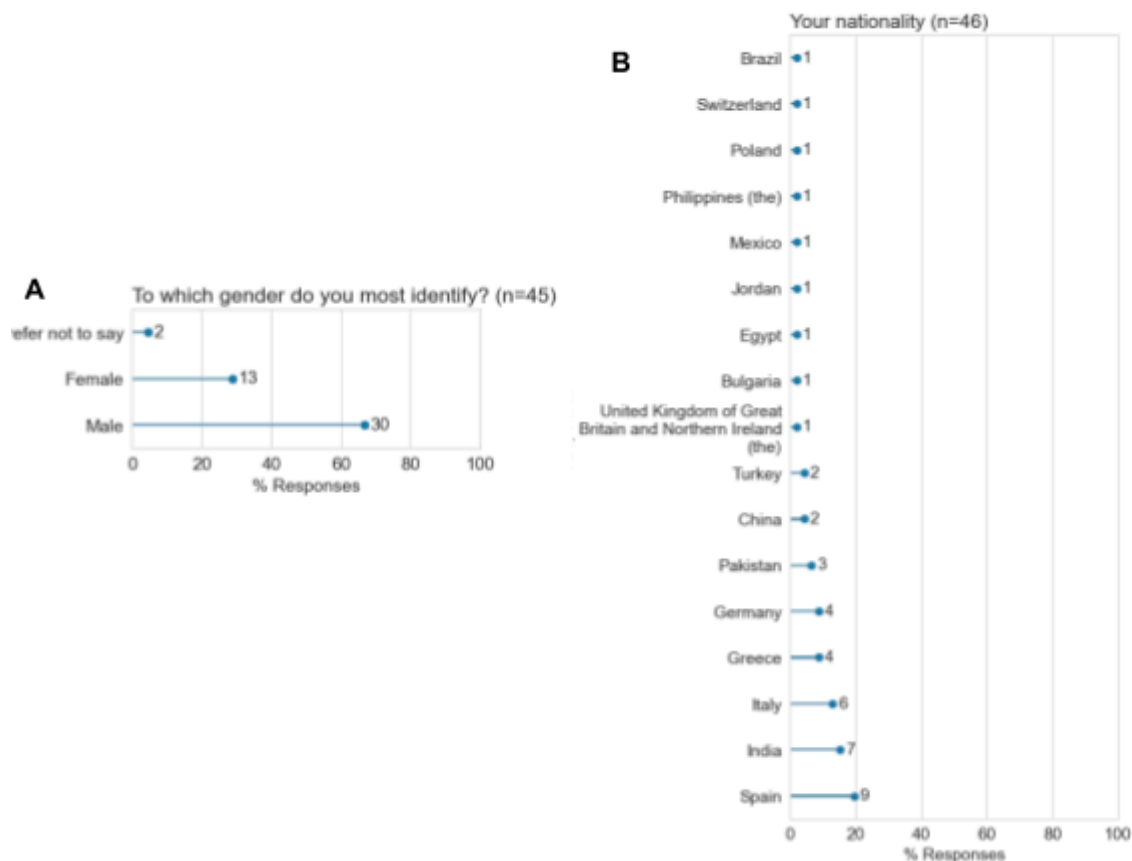


Figure 3 - Workshop applicants' nationality (A) and country of employment (B) from the pre-workshop survey.

As previously described in this report, participants were pre-selected primarily based on their country of employment, prioritising those based in Europe. Most of the selected participants were from Spain, followed by notable groups from Italy, Germany and Greece.

Interestingly, while employment was largely Europe-focused, participant nationalities were globally diverse. All participants came from countries across Europe, Asia, and Africa, reflecting a rich international mix within a predominantly European employment context. This distribution highlights the targeted geographic selection while still embracing a broad spectrum of global talent.

The survey participants represent a highly specialised group, primarily concentrated in the field of High Performance Computing and Computer Sciences already (~60%) (Figure 4). Nearly half of respondents reported working in life, biomedical or health sciences (~52%), being a well-represented group.




Figure 4 - The relevant fields or disciplines (multiple choice) question pre-workshop survey (n = 45).

The survey participants were predominantly experienced technical people (Figure 5A), with Research Staff (~31%) and Technical support staff (~26%) comprising the majority of the respondents. Other technical positions dominate the rest of the group, including industry employees, analysts, researchers in various career stages, and students recently graduated or undergrad. Despite their current employment, participants expressed a strong interest in transitioning to industry careers, particularly within SMEs (Figure 5B). Nearly all participants reported some connection to an SME. Specifically, 9 participants currently work in SMEs, 11 already collaborate with one, and a significant number (n = 11) indicated a desire to work in an SME. This contrast between current academic employment and SME connection suggests that the survey captured researchers collaborating with SMEs or actively seeking to transition from academia to industry, or for students to have better employability skills.

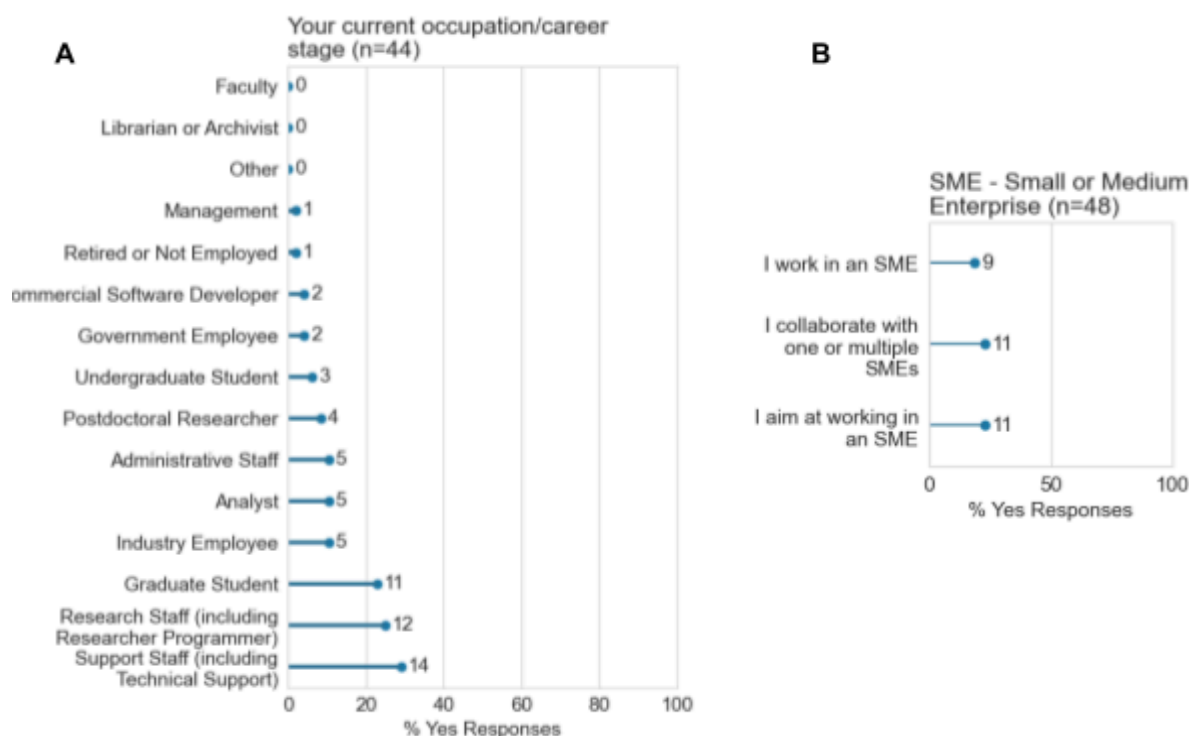


Figure 5 - The current occupation/career stage (multiple choice; n = 44 respondents) (**A**); Connection of the workshop participants to SMEs (n = 48 respondents) (**B**) according to the pre-workshop survey.

Applicants found information about the workshop through various channels, as illustrated in Figure 6, with the majority learning about it via social media or receiving direct recommendations from friends or colleagues.

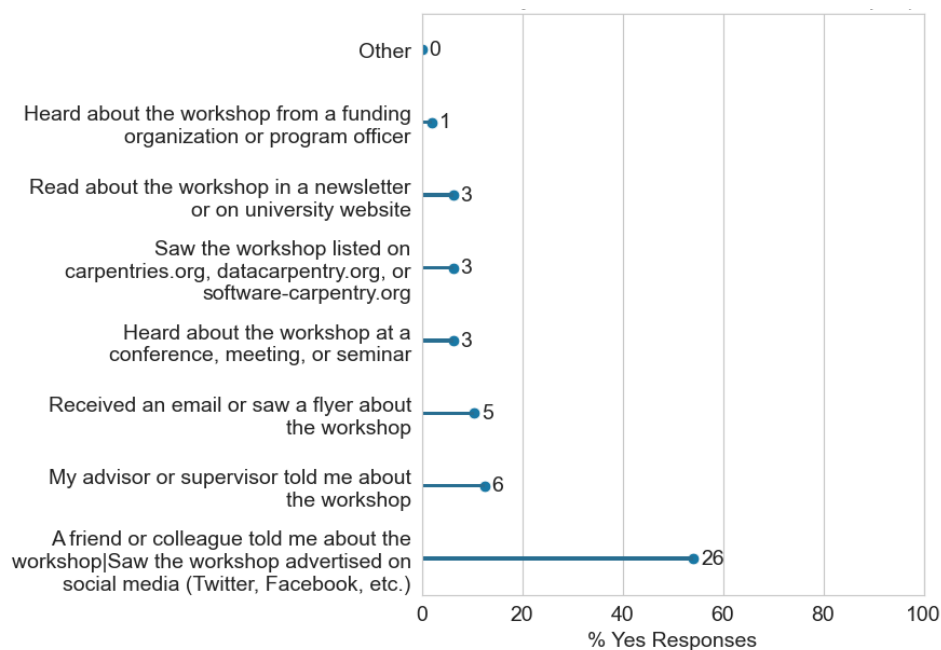


Figure 6 - Answers to the question: "How did you find out about this workshop? (multiple choice)" (n = 42).

The survey data reveals a distinct divide between participants' comfort with command-line and programming tools versus traditional GUI-based specialised software. A majority of participants demonstrate advanced computational skills, with ~58% using programming languages (R, Python, C++, etc.) daily and an additional ~23% using them weekly (Figure 7A). This high percentage of programming language usage (total of daily and weekly usage) aligns with the ~82% who use command shell interfaces daily or weekly, indicating a participant base that is highly comfortable with text-based, technical computing environments (Figure 7B). The strong adoption of version control systems further reinforces this technical specialisation, with 47% using Git or similar tools daily, suggesting collaborative coding practices and systematic project management approaches (Figure 7C).

In contrast, the usage patterns for GUI-based specialised software and database systems show much more varied adoption levels. The 35% of participants use specialised software like SPSS, SAS, or GIS tools on a weekly or daily basis, the rest of usages are quite scattered and even (Figure 7D). Database usage also represents usages of all kinds, with ~47% using SQL or similar systems monthly or more frequently, while ~13% never use databases at all (Figure 7E). This pattern confirms that participants are highly skilled technicians in the High Performance Computing world, where GUI applications and Databases are not used in this kind of environment.

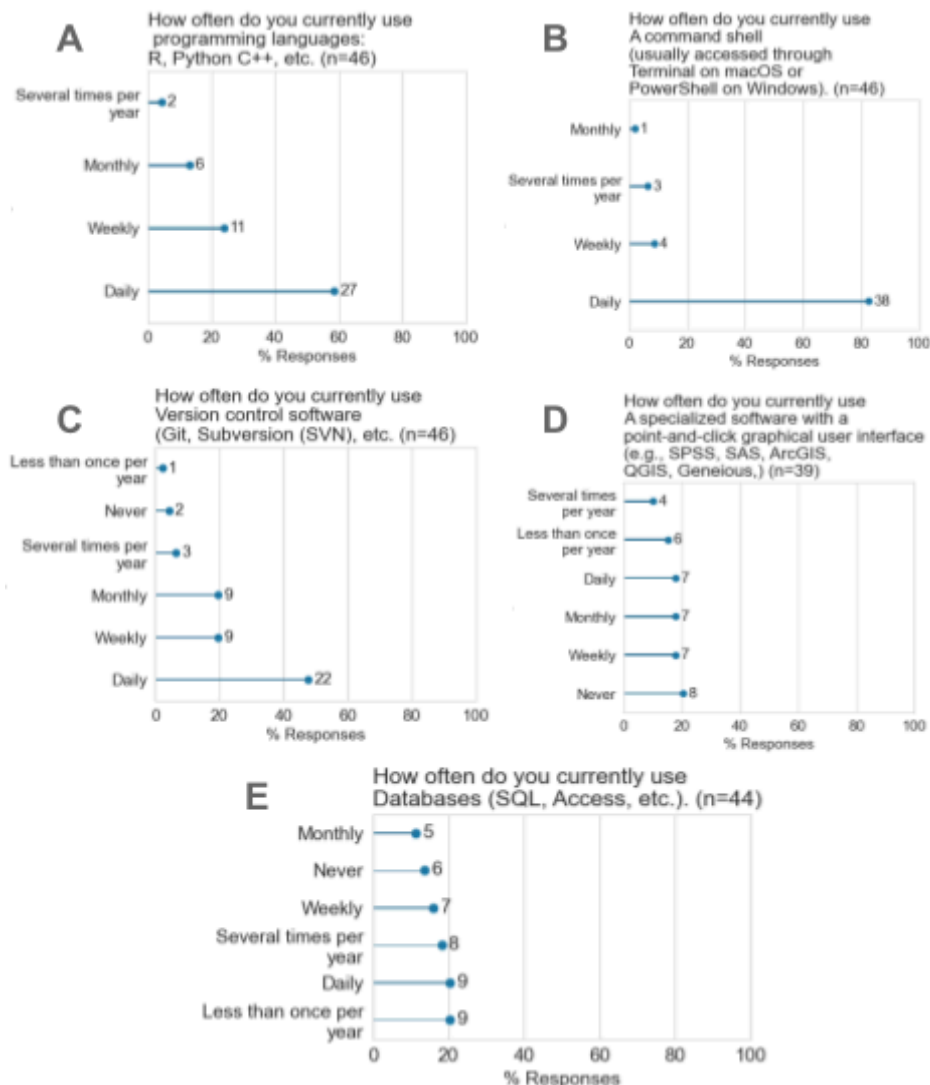


Figure 7 - Answers to the question: “How often do you use any of the following?”. This question was optional and trainees could answer none, some, or all the questions (n = 39-46): **A** - Programming languages: R, Python, C++, etc.; **B** - Terminal and macOS or PowerShell on Windows; **C** - Version control software (Git, Subversion (SVN), etc.); **D** - A specialised software with a point-and-click graphical user interface (e.g. SPSS, SAS, ArcGIS, QGIS, Geneious); **E** - Databases (SQL, Access, etc.).

Most participants were split between acquiring new skills (~75%), with some specifically interested in learning those applicable to their future job (~75%) and current occupation (~68%). Interestingly, about 31% of the participants expressed their intentions to leverage the acquired skills either to secure a promotion within their current job or to pursue new employment opportunities (n = 15) (Figure 8).

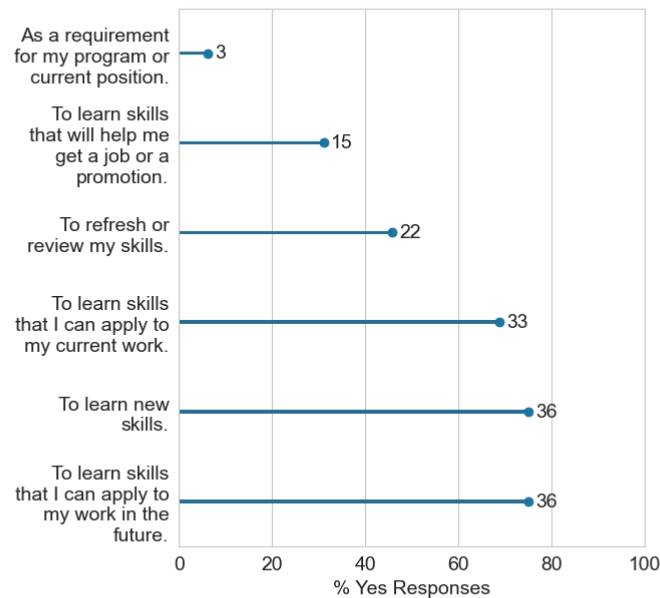


Figure 8 - Answers to the question: “Why are you participating in this workshop? (multiple choice)” (n = 48).

Participation

All 30 selected applicants of Module 2 attended the workshop live, although not all of them stayed for the entire duration (Table 3). All the selected participants received the self-learning materials for them to consult at any time.

Day	Participants	Instructors	Helpers
1	33	3	5
2	32	2	5
3	30	3	5

Table 3 - Number of participants, instructors and helpers per day. The number of participants was obtained from the participant login information captured by Zoom.

Evaluation

19 workshop participants submitted the certificate request. Submissions were marked against a set of assessment criteria, needing to fulfil the following requirements, in each trainee VM:

Requirements criteria

- Operational users created correctly
- Have Warewolf installed
- Working “sinfo” command
- Monitoring dockers up and running

- Ansible installed and correct playbooks
- “Module avail” command working

In summary, all 19 participants who submitted the request fulfilled the criteria; therefore, all of them received a certificate (Figure 9).



Figure 9 - Certificate provided to attendees who successfully completed the assignments

Impact and outcomes

Daily feedback

At the end of each day, participants were asked for feedback on the following three points:

- Please share one thing that was good about today
- Please share one thing that could be improved about today
- Do you have any other comments?

The workshop garnered positive feedback, especially for its well-defined structure, organisation, and content. More specifically, the initial day received mixed reviews; most participants thanked the instructors for the great sessions, but felt that the hands-on part was confusing to follow and that it was too much copy-paste sometimes. With these reviews, the content of the second and third days were adjusted so that participants would be less prone to errors in the hands-on, and they could focus more on the explanation of the trainers. The feedback from these days reflected the positive changes to the training content.

Post-workshop survey

At the end of the workshop, participants were invited to complete a post-workshop survey consisting of 15 questions, some of which were optional. A total of 24 participants completed the survey.

The evaluation data reveal high satisfaction levels across all measured dimensions, indicating a successful educational experience. All participants responded with "Strongly Agree" or "Agree" that they felt comfortable with the workshop environment (Figure 10A), creating the foundation for effective learning. This positive atmosphere translated into strong perceived practical value, with 73% of participants confident they can immediately apply the learned material to their work or research (i.e., 17 out of 23 responding with "Agree" and "Strongly Agree"; Figure 10B). The immediate applicability rating is particularly significant given the technical nature of System Administration concepts, with participants now being able to introduce concepts taught in the workshop to their daily work.

Instructor effectiveness emerged as a standout strength of the program, receiving consistently positive ratings across all evaluated criteria. For example, survey questions assessing instructors' clarity of answers (Figure 10C), enthusiasm (Figure 10D), and approachability (Figure 10E) showed that ~91% of participants (21 out of 24) responded positively with "Strongly Agree" or "Agree" answers, and at worst, a couple of "Neutral" answers. The survey also evaluated the instructors' knowledge by the scoring of the statement "The instructors were knowledgeable about the materials being taught". Here, 100% of the participants selected "Strongly Agree" or "Agree", reflecting a high level of confidence in the instructor's expertise (Figure 10F).

Notably, one participant consistently provided negative feedback across most categories, suggesting an outlier experience, while the remaining responses demonstrated consistency in positive evaluation. These results indicate that the workshop successfully delivered high-quality instruction in a supportive learning environment with strong practical relevance for participants.

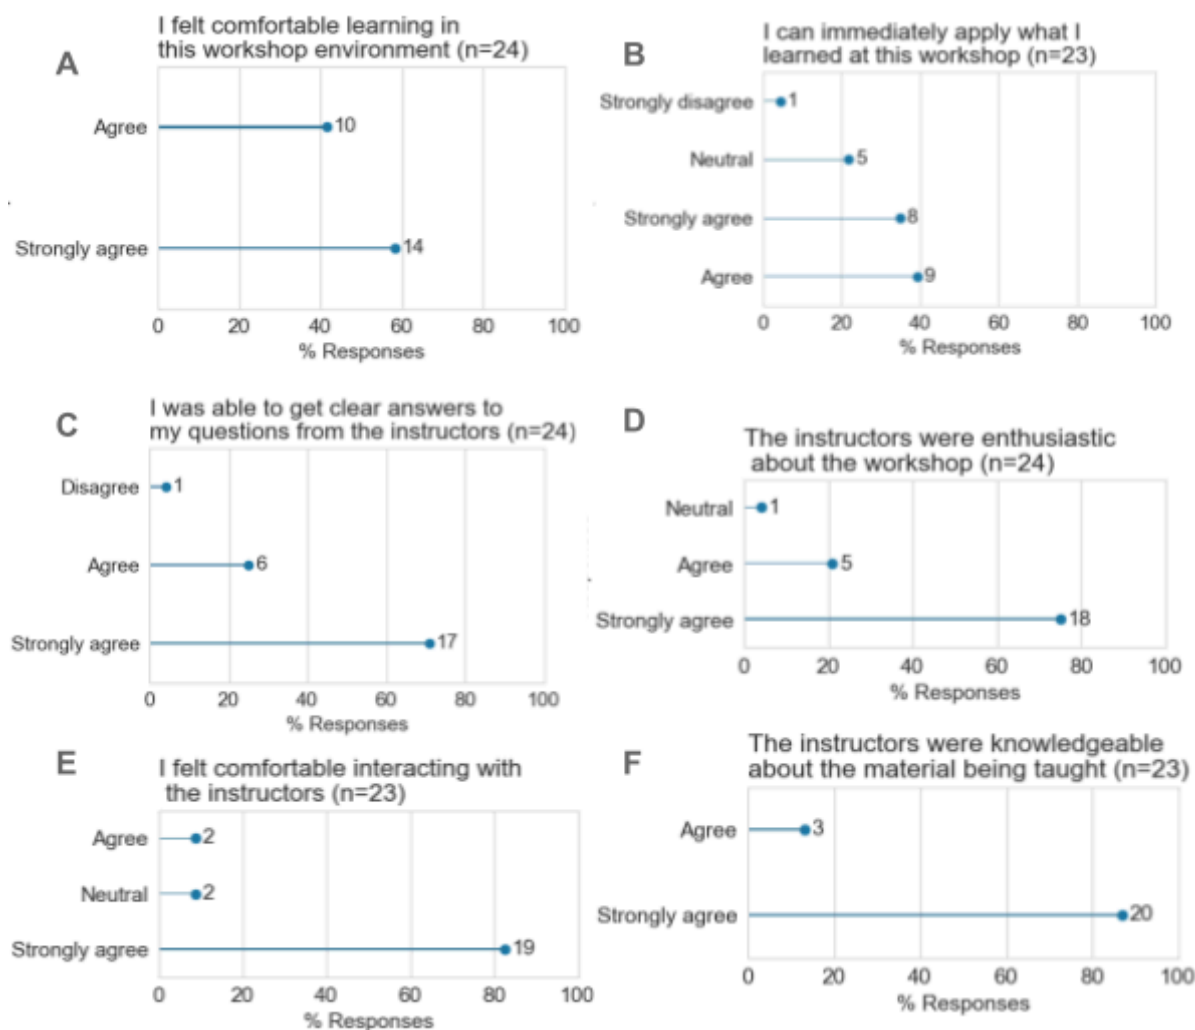


Figure 10 - Rating of participant's agreement with the following statements: **A** - I felt comfortable learning in this workshop environment; **B** - I can immediately apply what I learned at this workshop (n = 24); **C** - The instructors were enthusiastic about the workshop (n = 24); **D** - The instructors were enthusiastic about the workshop (n = 24); **E** - I felt comfortable interacting with the instructors (n = 23); **F** - The instructors were knowledgeable about the material being taught (n = 23).

This instructor's excellence, coupled with the optimum learning environment, culminated in high recommendation likelihood scores (Figure 11), with ~95% providing ratings of 8 or higher (respondents who chose "8", "9" and "10" were 4, 4 and 14 correspondingly). This suggests the workshop achieved its educational objectives for the vast majority of attendees, creating a replicable model for future system administration education.




Figure 11 - Answers to the question: How likely would participants recommend this workshop to a friend or colleague? (n = 24); Rating interpretation - 1-3: No, 4-7: Maybe, 8-10: Yes; The answers from 0 to 6 are not shown in the plot as no responses were given in this range.

Overall feedback

The post-workshop survey assessed participants' feedback on their learning experience using two statements: "Please list the major strengths of this workshop", and "Please list the ways the workshop could be improved". Recurring themes from the responses to these questions were identified; all responses are about the same 4 pillars. This indicates that participants found consistent value in the overall workshop experience.

Instructor quality, not only in explanations but also when troubleshooting issues, highlighted the impact of knowledgeable, experienced instructors who can effectively bridge theoretical concepts with real-world applications. Content Quality, especially the hands-on sessions, emerged as one of the program's greatest strengths, being mentioned by the trainees as excellent for deeply understanding the theoretical content and how to apply it. Related to this, the feedback commends the technical infrastructure that supported the course, as each participant could install and deploy their own mini HPC cluster.

The improvement suggestions (i.e., ways the workshop could be improved), rather than indicating fundamental flaws, pointed toward opportunities to enhance an already solid educational foundation. All messages also revolve around three clear topics, and interestingly enough, three spare trainees state "nothing to be improved", and a couple of spare points just complain about the coffee breaks and lunch time, as this was held in a southern European schedule.

The first improvement area is fewer dense theoretical sessions, as sometimes it was long explanations, and they prefer to just go to the core theoretical explanation and then complement it with the hands-on part. The other big suggestion is a better explanation of the issues, as, while instructors could solve all of the trainees' problems, sometimes there was not enough feedback on what the actual error was. The remaining improvement area relates to the first one, and is to give more time for theoretical and hands-on sessions, as they were

very condensed. This pattern suggests that participants were really happy about the training and were not seeking any big change, but rather incremental improvements that would elevate their learning experience from very good to exceptional, indicating a mature program ready for targeted enhancements rather than complete revision.

Participants' self-reported improvement in skill and confidence

The pre- and post-workshop assessment revealed consistent and substantial improvements. The most significant improvements occurred in practical system administration applications. Trainees praised in the feedback section their new ability to build their own HPC cluster, although a tiny one; they seemed very enthusiastic about it. These gains are especially meaningful given that participants already possessed considerable technical backgrounds, suggesting the workshop effectively advanced existing competencies rather than merely introducing basic concepts. These overall improvements in skill and confidence levels indicate that the workshop ultimately positioned participants for better management of High Performance Computing machines.

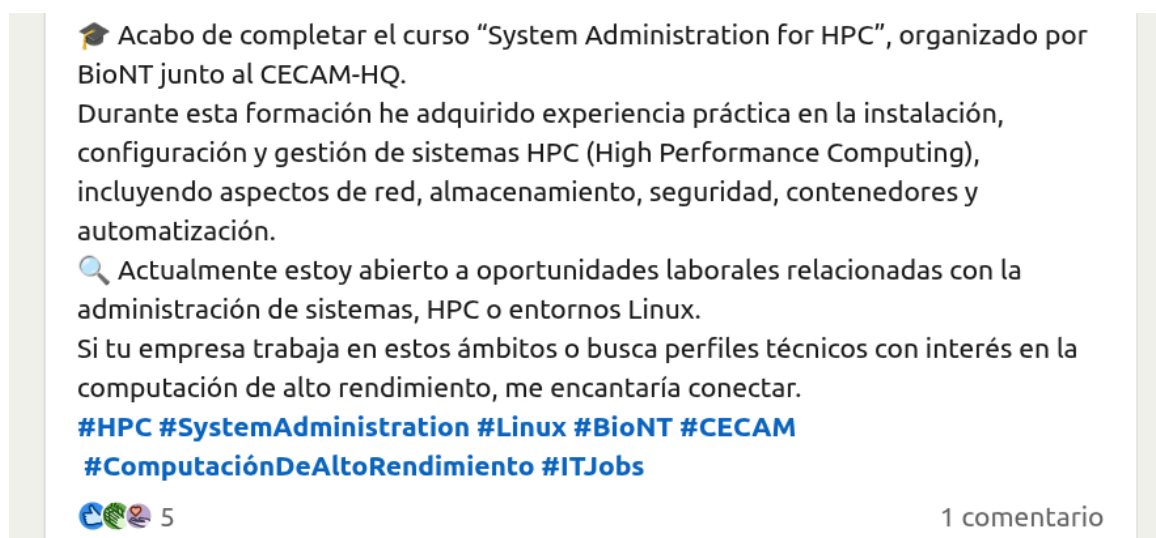


Figure 12 - Public post of a participant on their LinkedIn profile endorsing the workshop (names of the participants were not shown to maintain anonymity). Translation to English: "I have just completed the course "System Administration for HPC," organized by BioNT in collaboration with CECAM-HQ. During this training, I gained practical experience in the installation, configuration, and management of HPC (High Performance Computing) systems, including networking, storage, security, containers, and automation. I am currently open to job opportunities related to system administration, HPC, or Linux environments. If your company works in these areas or is looking for technical profiles with an interest in high-performance computing, I would love to connect."

Conclusion

The 8th BioNT workshop, "System administration for HPC," successfully took place between October 21st and 23rd, 2025. To complement the main workshop, an industry meetup event titled "Sysadmins in HPC meet-up" was organised on October 10th. The event was designed as a hybrid format with live streaming to attract a wider audience. This meetup event

effectively promoted the workshop to industry employees, students and academics interested in transitioning to industry roles, particularly within SMEs. The success of this outreach strategy is evidenced by the survey responses discussed earlier in this report.

The setup for the 9th workshop benefited from the experience gathered during the first eight BioNT workshops and had additional attention to actively reach the audience, making sure the audience had a chance to fulfil the prerequisites.

This workshop saw a good engagement from pre-workshop events and through registration. There was an oversubscription, and 45 participants were selected using a consistent criterion to project goals, where participants in Europe who were seeking a job or were involved in an SME received priority.

The heightened interest, as evidenced by good participation numbers, can be attributed to more extensive advertising efforts, highlighting the effectiveness of increased partner engagement in promotions. One of the distinguishing factors for the main workshop was its invited instructors, who are professionals in the field. Another important feature of the workshop was that it included real-world case examples of extensive hands-on sessions on the topic.

There were some significant challenges during this workshop, mostly related to the planning of the specialised infrastructure needed for the course, the preparation of the VM's over AWS infrastructure faced technical difficulties and administrative ones, as a quota increase had to be requested to allocate all the VM's, and afterwards some AWS regions didn't have the capacity to spin all the instances. This was solved by changing the instance type; although they were more expensive, it was still under the allocated budget in the project.

Despite the mentioned risks and challenges, BioNT partners contributed heavily to this workshop to make it a success. As discussed in various sections of the report, participants consistently reported a comfortable and supportive learning environment, high satisfaction with the knowledgeable and enthusiastic instructors, and the immediate applicability of the course material. The workshop boosted participants' self-reported skills and confidence in several areas. The high positive recommendation rate, coupled with feedback received during and after the event, further confirms the workshop's success in meeting participant needs and expectations.