

One Year HPC Certification Forum in Retrospective

Julian Kunkel
University of Reading
Reading, United Kingdom
j.m.kunkel@reading.ac.uk

Kai Himstedt
Universität Hamburg
Hamburg, Germany

Weronika Flinger
EPCC, The University of Edinburgh
Edinburgh, United Kingdom

Jean-Thomas Acquaviva
DDN
Paris, France

Anja Gerbes
Goethe-Universität
Frankfurt am Main, Germany

Lev Lafayette
University of Melbourne
Melburne, Australia

ABSTRACT

The ever-changing nature of HPC has always compelled the HPC community to focus a lot of effort into training of new and existing practitioners. Historically, these efforts were tailored around a typical group of users possessing, due to their background, a certain set of programming skills. However, as HPC has become more diverse in terms of hardware, software and the user background, the traditional training approaches became insufficient in addressing training needs of our community. This increasingly complicated HPC landscape makes development and delivery of new training materials challenging. How should we develop training for users, often coming from non-traditionally HPC disciplines, and only interested in learning a particular set of skills? How can we satisfy their training needs if we don't really understand what these are? It's clear that HPC centres struggle to identify and overcome the gaps in users' knowledge, while users struggle to identify skills required to perform their tasks.

With the HPC Certification Forum, we aim to clearly categorise, define, and examine competencies expected from proficient HPC practitioners. In this article, we report the status and progress this independent body has made during the first year of its existence. The drafted processes and prototypes are expected to mature into a holistic ecosystem beneficial for all stakeholders in HPC education.

1 INTRODUCTION

There is a generally accepted set of skills and competencies necessary to efficiently use HPC resources. This skill set depends on the role and domain of the practitioner but also on the available infrastructure of the centre providing the computing resources. For example, a scientist needing to run an application on a specific machine may need basic skills in Linux, MPI, environment modules, and knowledge about the batch scheduler, e.g., Slurm. Now rather than providing that scientists with a list of instructions to be followed in a copy and paste fashion or a never-ending list of things they should know, we want them to understand each of the required steps without having to learn everything about it, but at the same time understand how it fits into the bigger picture of

using an HPC system. For instance, understanding Slurm is a good example of a very fine-grained skill that fits under a more generic skill defined as "resource management", illustrating concepts across the rich variety of available resource managers.

Most data centers operating HPC systems offer regular training events focusing on general aspects of their supercomputer's hardware architecture, software stack, application development environment and various tuning and debugging tools. It is understandable why the materials they offer are geared towards the special demands of the specific HPC environment and the institutions they support. The problem is that teaching content typically covers only a fraction of the HPC skills necessary to use another HPC system. This approach of teaching only specific implementations, tools or workflows—instead of concepts behind them—makes a move to a different system or tool unnecessary complicated. Narrowing the scope of training events make sense from the HPC provider perspective, but is not very conducive to development of a more comprehensive learning environment that provides assessment and certifies the newly acquired skills. Although certificates are used widely in IT industry to verify certain knowledge, until now there was no similar approach for HPC training. It is, however, clear that a mapping of competences and certification scheme could address some of the challenges stemming from the constantly growing training needs of our community,

This article describes the current status of the certification program curated by the HPC Certification Forum. Our previous work provided a brief overview of the evolution from the project that sparked the HPC Certification Forum [2] while this article provides details on adopted mechanisms and design decisions.

The article is structured as follows: First, in Section 2 we introduce the HPC Certification Forum and the certification program. Then, an overview of the status is given in Section 3, followed by the organisation of competencies discussed in Section 4. This leads to the proposal for the certification process in Section 5. In Section 6 the whole learning ecosystem is described. Related work is presented in Section 7. Finally, the article is concluded in Section 8.

2 THE HPC CERTIFICATION FORUM

The HPC Certification Forum (HPCCF) has the role of a (virtual) central authority to curate and maintain the proposed certification program. The program consists of three parts: the tree of defined competencies, the examination of practitioners to prove they possess those skills, and finally the certification demonstrating their knowledge. Although, the forum is not involved in development of any

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Copyright ©JOCSE, a supported publication of the Shodor Education Foundation Inc.

training materials or tools, it supports the ecosystem around the competencies.

The HPCCF aims to support existing activities and complements them by providing a unified and clear way of mapping out the relevant HPC competencies. Thus, the HPCCF does not regulate the content of training material; we purposely separate the definition of skills, the examination and the certification from content delivery. Similarly, the program does not prescribe a curriculum or any fixed order by which skills need to be obtained. It eases the navigation between different competencies without being overly restrictive.

The forum is organised around a Webpage¹, a GitHub repository², and a team collaboration tool (Slack) for communication and monthly meetings. The membership is divided into three categories: associate, full and steering board. Anyone can become an associate member free of charge, allowing a passive involvement – e.g., observing the status, being listed on the webpage. Full members must actively contribute to the program and, in return, gain voting rights for the roles of the steering board. The steering board directs the overall activities and is organised in different responsibilities, i.e., into topic-specific chairs. At the moment the steering board has 9 members.

3 OVERVIEW OF THE HPCCF ACTIVITIES

The first year of the HPCCF forum was rich in discussions, both internally and with external partners, resulting in a number of processes and prototype tools being put in place. These are described briefly below, and in more details in the subsequent sections.

Management. The first steering board has been elected at ISC-HPC 2018. We established communication channels via webpage, email and Slack for regular discussion. Slack is also used to conduct our monthly open meetings between members and the steering board. At beginning, video conferencing tools have been used but the geographical diversity of the attendees influenced the decision to move towards chat-based tools. The asynchronous nature and automatic message retention proved to be more productive, and provided a more open and inclusive platform for discussion. Meeting notes with action items are extracted from Slack and documented using Google Doc (an online word processor allowing real-time and multiple-user collaboration).

The ongoing conversation with various stakeholders involved in HPC education and training, keeps producing more engagement in the forum's activities. The member's contributions are voluntarily and follow the schedule agreed between the contributor and the HPCCF board. The prospective contributions are managed in Trello (a web-based list-making collaborative application), in which each member can manage their own card.

Technical tasks. The tasks we set out to accomplish are: 1) the definition of competencies; 2) the examination of practitioners; 3) the creation of certificates; and 4) the reinforcement of an ecosystem of tools supporting them.

For all these sub-goals we made substantial progress. While the definition and organisation of competencies was the main focus,

we prototyped various tools and processes that embed the competencies into the wider education ecosystem.

Firstly, considering the pedagogic literature from higher education we finalised the templates for the skills description. The effort to map out all of the relevant competencies is still on-going, but the feedback received during a number of events organised by the HPCCF and other members of the HPC education community resulted in the addition of new topics to the first version of the certification program.

The definitions of the competencies are version controlled with Git, and are publicly available in XML and Markdown format. A Wiki³ is being used to edit the skills' descriptions, and a navigable and widely customisable JavaScript⁴ makes the skill tree easily accessible to the forum members and the end-users alike.

Preliminary processes and tools have been developed and deployed for curation of the examination questions and for conducting an online multiple-choice examination.

4 SKILLS

A skill is defined as a set of learning outcomes and relevant metadata. Within a single skill, there can also be multiple levels (basic, intermediate and expert level) building upon each other and further distinguishing the expertise. We expect the practitioners to acquire the lower levels before progressing to more complex levels.

The basic level should cover the most relevant aspect of the skill (needed by anyone that uses the skill), whereas the intermediate (used for common exceptional cases) and expert levels (used in special circumstances) are needed only for a subset of users. Typically, an institution has only few experts, if any at all.

This model can be compared to the classification of school knowledge, for example, a skill with the name "basic arithmetic operations" would include the math skills of addition (being able to add numbers) and subtraction (being able to subtract numbers), multiplication (being able to multiply numbers) and division (being able to divide numbers). Operating on two numbers with \pm would be a basic level of the skill, whereas including multiplication/division and mentally operating on hundreds of numbers would be an expert level. Another example could be a skill called "technical drawing", in which being able to draw simple geometric figures would be counted as a basic level, and an ability to draw a complicated mechanical objects would be an expert level.

The skills are organised in a tree structure from a coarse-grained to a fine-grained representation, allowing users to browse the skill based on the semantics. The root level of our current⁵ skill tree is shown in Figure 1. The basic idea is that the skills from the root level simplify the navigation by providing an indication about their scope, e.g., core knowledge, usage of HPC environments, or about programming. This should allow the user to rapidly drill into the skill representation. The closer a skill is to the root, the more abstractly it is defined, while leaf nodes cover the knowledge for a specific skill.

As the tree serves the purpose of organising the skills, the references from one branch to a skill in another branch are allowed. This

¹<http://hpc-certification.org>

²<https://github.com/HPC-certification-forum>

³<https://www.hpc-certification.org/wiki/>

⁴<https://www.hpc-certification.org/skills/map/>

⁵The goal is to finalise the tree this year.

important feature makes the reuse of the skill definitions possible, at the same time allowing users to navigate the tree according to the semantics.

4.1 Description of a skill

Each skill on the tree, including the inner nodes, is described in more detail as follows:

- **ID:** Identifier according to its position in the skill tree. The last character indicates the level of the skill (Basic, Intermediate, or Advanced).
- **Name:** A name capturing the essence of the skill.
- **Background:** Provides brief information motivating the need for the skill and how it fits into the bigger picture with other skills.
- **Aims**⁶: Describe the purpose of the skills, but doesn't really include a list of what a practitioner will learn or do. Explaining what a skill is trying to achieve is not the same as saying how it should be done.
- **Learning outcomes (LOs):** Defines briefly what practitioners will learn. The objectives are statements what prospective learners are able to do. They should clearly describe or define an action bringing about a measurable/quantifiable increase in understanding of that skill.

On the leaf level, a skill is fine-grained and orthogonal to other skills – their narrowed scope means they can be taught in sessions ranging from a 1.5 hour lecture up to a 4 hour workshop. We believe this granularity allows practitioners to cherry-pick the skills relevant to their circumstances, and lecturers and examiners to prepare small lectures with well-defined content. For technology-dependent skills on the leaf level (e.g., a specific file system or workload manager) the introductory skills are often provided, as they contribute to the foundation of many specialised skills representing a specific hardware or software technology.

The aggregation within the tree is similar to that found in the education circles. Hussey et al.[1] categorised the aggregation of skills into the following levels: individual teaching events, specified for modules or short courses, and those specified for whole degree programmes. Hussey et al. conclude that LOs on the coarse grained level are not useful for high-level education because they “would state little more than an annotated list of contents” and relevant for those familiar with the subject. To overcome these issues, we provide a relaxed level of granularity and lower level of abstraction of the learning outcomes on inner nodes.

4.2 Skill Examples

This subsection presents an example of an inner-node skill called "Executing parallel applications", as well as its two leaf skills - "Workload manager introduction" and "Slurm workload manager".

- **ID:** USE4.2-B
- **Name:** Executing parallel applications
- **Background:** Parallel computers are operated differently than a normal PC – all users share the system. To ensure

a fair and efficient use of shared resources, various operative procedures are put in place. Users must understand these concepts and procedures to be able to run a parallel application on the resources available to them. Moreover, different HPC systems adopt different solution to manage their resources.

- **Aim:** To enable practitioners to comprehend the concepts and procedures for running parallel applications in HPC environments; to run and monitor the execution of parallel applications on HPC systems.
- **Learning outcomes:**
 - explain the concepts and procedures related to resource allocation and job execution in an HPC environment;
 - run interactive jobs and batch jobs;
 - understand and describe the content and expected behaviour of job scripts;
 - change provided job scripts and embed them into shell scripts to run a variety of parallel applications;
 - analyse the output generated from a job scheduler and understand the cause of typically generated errors.

The next leaf-level skill describes how workload management works in general, regardless of the specific software implementing it. The aim and outcomes defined by the parent skill (described above) are expected to be covered to a certain extend and refined by this leaf skill.

- **ID:** USE4.2.1-B
- **Name:** Workload manager introduction
- **Background:** There is a wide range of different workload managers in use. This skill on a conceptual level explains how to use them.
- **Aim:** To enable practitioners to comprehend and describe the basic architecture and concepts of resource allocation on an HPC system.
- **Learning outcomes:**
 - comprehend the exclusive and shared usage model in HPC;
 - explain the generic steps required to run and monitor a single job;
 - differentiate between the batch and interactive job submission;
 - comprehend the generic concepts and structure of resource manager, scheduler, job and job script;
 - explain the role of environment variables as a mean to communicate certain settings of your job;
 - comprehend job budgeting and accounting principles.

The following skill describes a leaf-level skill for the usage of Slurm at the basic level. While there is no formal dependency to the introduction (skill USE4.2.2-B), it is expected that practitioners have obtained the other skill before learning this one.

- **ID:** USE4.2.2-B
- **Name:** Slurm Workload manager
- **Background:** Slurm is a widely used open-source workload manager, which also provides various advanced features.
- **Aims:**
 - To enable practitioners to use relevant tools to run and monitor parallel applications using Slurm.

⁶The definition of aims and outcomes follows literature for higher education [5] and <https://www.heacademy.ac.uk/system/files/assessment-learning-outcomes.pdf>.

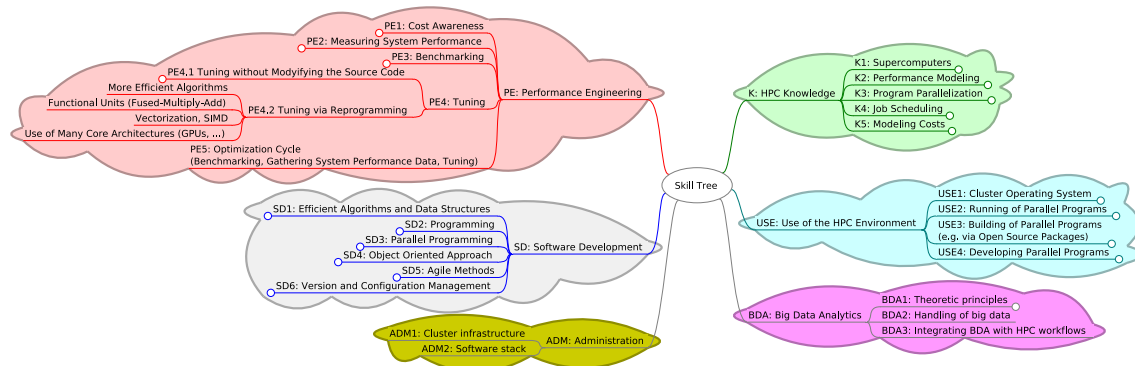


Figure 1: Skill tree: top level competencies with the PE4 branch expanded.

- To enable practitioners to comprehend and describe the basic structure of Slurm and the related suite of tools.
- **Learning outcomes:**
 - run interactive jobs with *salloc*, a batch job with *sbatch*;
 - explain the architecture of Slurm, i.e., the role of *Slurmd*, *srun* and the injection of environment variables;
 - explain the function of the tools: *sacct*, *sbatch*, *salloc*, *srun*, *scancel*, *squeue* and *sinfo*;
 - explain time limits and the benefit of a backfill scheduler;
 - comprehend that environment variables are set when running a job;
 - comprehend and describe the expected behaviour of a simple job script;
 - comprehend how variables are prioritised when using command line and a script;
 - change a provided job template and embed them into shell scripts to run a variety of parallel applications;
 - analyse the output generated from submitting to the job scheduler and typically generated errors.

The learning outcomes of this skill, marked with B to indicate its basic level, are focusing on the most fundamental knowledge required to use Slurm effectively. An intermediate version of the skill could, for example, cover how to create reservations or special queues.

5 CERTIFICATION

A certificate shall serve as a confirmation that a user obtained the expertise in the relevant skills. Special considerations needs to be given to the certification integrity. We must ensure that the learning outcomes of the covered skills are examined to a satisfactory degree, and prevent self-cheating or dishonesty. Students can always cheat during an examination – the susceptibility to cheating depends on many factors, including but not limited to the type of exam, methods used for assessment, motivation for taking it, consequences of failing or scoring poorly, and the possibility and frequency of resits.

We want to enable a cost-effective assessment (free for the practitioner), so we chose an online examination. This increases the opportunity for cheating [3] but we plan to deploy a several strategies to minimise the risk of cheating, such as raising the awareness of the examinees, using a pool of questions, time limits for each

question and a delay between registering for and taking the examination.

Since knowledge can age, each certificate needs to indicated when (month and year) the qualifying examination took place. Also, because the examination of a single fine-grained leaf-level skill would be too easy for short-term memorisation and more prone to self-cheating, the certificates bundle multiple skills together. We believe the incentive to deliberately fake the assessment, e.g., by having the exam filled by someone else, and thus, being awarded a certificate, is low. Therefore, we address this issue in a lightweight fashion only.

A process has also been created for prospective contributors of the examination questions. The questions are the only proprietary component for the HPCCF – using restrictive license terms for authors while giving them credit. This is considered necessary for managing the database containing solutions to the examination questions.

In the next section, we discuss the process proposed for conducting the examination.

5.1 Examination Process

The examination process can be started at any time but a delay is employed before the examination actually starts. Firstly, a user has to register using their name, email address and optionally an affiliation. This process is explained, together with privacy policies, on the examination website. Next, they need to read a paragraph on the integrity policy defining cheating and encouraging honesty. Practitioners are asked to opt-in for the policy statement. Upon receiving this registration, we generate an encrypted token on the server that is returned to the user by email. Up to this stage, we do not store any information about the user on the server.

The email contains a link that will start the examination, and the user chooses to start the examination by clicking on it. Only then, the information about the user and the examination start time is stored in a temporary database.

The summative assessment itself is conducted using multiple choice questions. Therefore, for each skill, we will develop a pool of questions and answers, and accept external contributions. Each question has a pool of possible answers (e.g., 10). An examination

consists of randomly selected questions with five of their answers, and the practitioner has to decide if each statement is true or false.

Once the user submits the completed exam, the selection of answers is stored on the server for validation. At the moment, we envision the validation process includes a final manual step to trigger the certificate generation once provided with all information. If the user meets the pass criteria (typically 70% of correct answers) the earned certificate is issued and sent to the user by email. All personal information about the user is then deleted from the server, but the affiliation and raw responses are preserved. The affiliation is used for promotional purposes, while the raw responses will be analysed to optimise the questions. For example, if we identify that most users make the same mistake it may be an indication that either that question or its answer is too ambiguous and should be improved.

If the user didn't meet the pass criteria, they will be informed about their score. The user can then retry to obtain the certificate after a cool-down period (typically one week), but not immediately afterwards to prevent success via brute force methods. To enforce this, the information about when the exam was undertaken is linked to the user's name and email. As the sole purpose of the multiple choice questions is the examination, the incorrectly answered questions will not be revealed. We assume that high quality training materials available within the HPC community, and covering the individual skills being part of this certification program, offer constructive feedback and other ways of consolidating the newly acquired knowledge. Therefore, the Forum's only concern is to provide the mechanism of certifying whether the learners possess that knowledge or not; As opposed to pointing out gaps in their understanding.

5.2 Certificates

The examinees that pass the examination are awarded a corresponding certificate. Such certificate consists of two parts: a PDF and a text file. The PDF contains the key information making the certificate meaningful. An example of how it may look like is presented in Figure 2. The name and identifier of the certificate is found in the centre, on our example these are "HPC Driving License" and ID 1, respectively. Similar to a driving license, this particular certificate could provide the minimum set of knowledge required to understand and use a typical supercomputer.

The text file contains the same information, as well as a verification URL that can be given to a third-party to confirm the certificate's credibility. It is also PGP signed using the private key of the HPC Certification Forum to allow verification with the public key. An example file looks as follows:

```
-----BEGIN PGP SIGNED MESSAGE-----
Hash: SHA512
HPC Certification Forum Certificate
This text confirms that "Julian M. Kunkel" has
successfully obtained the certificate
"HPC driving license" (id: 1) at 02/2019.
Verification URL: https://hpc-certification.org/[...]
-----BEGIN PGP SIGNATURE-----
[...]
-----END PGP SIGNATURE-----
```



Figure 2: Draft for an awarded certificate

Together both documents should be able to provide enough credibility for the certification to be fully functional within the HPC academic and professional domains. Further improvements will be adopted as required.

6 ECOSYSTEM

The Forum encourages the development of an ecosystem around the HPC classification by supporting training and tool development. Many of its members are involved in those activities in a professional capacity, so the development of training materials and tools, and the development of the certification program have elements of co-design in them.

6.1 Training Delivery

The HPC Certification Forum is not developing training material directly or competing with providers of training material. However, we support individuals and institutions by endorsing and promoting their training materials and courses in two ways.

Firstly, an author is allowed to indicate on the training material itself or a related promotional material which skills are covered either fully or partially. We provide a seal that can be used for that purpose (see Figure 3). The reference to the HPCCF and the seal can be used free of charge under the condition that the developer of the training material registers a link to the material (or course) on our webpage using an online form. That way, the HPCCF is informed about the usage of the seal.

Secondly, we will link from our webpage the endorsed training material covering the individual skills and certificates. By using JavaScript and dynamic webpages, we will provide various views of the skills – with and without links to suitable training materials.

Note that we are not intending to verify the correct usage of the seal explicitly. However, in case the training material or course doesn't deliver the expected material practitioners may complain and we will remove the link to that training material from our webpage.

We expect that this strategy will make a good range of free training material available for most skills, while various institutions



This training covers (partially)
 - K1.1 System architectures
 - K1.2 Hardware architectures
 See <https://hpc-certification.org/c/1.0>

Figure 3: Draft for the seal for teaching material

and individuals can still charge for effective training courses. Ultimately, the catalogued training materials will complement each other, leading to a rich variety of content suitable for each individual practitioner, e.g., addressing different learning styles and languages. This becomes increasingly important as more in-time personalised training is needed by our ever-growing HPC community.

6.2 Navigation

Thanks to the contributions of the PeCoH project⁷, a JavaScript prototype embeds a navigable skill tree into a webpage. The script can be used by different stakeholders to adjust and present the skill tree in the most suitable for them view. By view we mean the selection and organisation of the skills, and the additional information provided when looking at the tree. For instance, you could have a tree for a specific scientific domain, for a specific data centre, or even for the role of the tester of a specific application. This feature could be used to indicate which skills are mandatory or beneficial for these use cases, providing links to training material or adding further supplementary information.

7 RELATED WORK

Relevant work can be divided into two categories: efforts to establish an HPC curriculum, and systematic efforts to design, develop and deliver HPC teaching and training materials.

In academia, many universities offer their own curriculum around scientific computing and HPC, as part of their undergraduate programmes, covering a range of theoretical and practical aspects relevant to the students field of study, e.g., software development of numerical domain-specific applications. Such courses rarely cover the efficient and effective use of HPC systems. The number of taught modules focusing on HPC is small, and the number of the whole programmes dedicated to HPC is even smaller. We are aware such

⁷<https://www.hcc.uni-hamburg.de/pecoh.html>

programmes are run by: EPCC at the University of Edinburgh⁸, the University of Liverpool⁹, the International School for Advanced Studies (SISSA) and ICTP in Italy¹⁰, the Trinity College Dublin¹¹ in Ireland, the Polytechnic University of Catalonia¹² and BSC in Spain and the University of Cote d'Azur¹³ in France. These 1-2 year master programmes are fully accredited by the universities that run them, but that also makes them fairly expensive and their comprehensiveness is not suitable for the majority of HPC users.

Data centres offer their own teaching materials and run a number of training courses to support their own users. Although, these are becoming more varied in content and more accessible through the adoption of online, remote and asynchronous delivery methods, they are still not enough to fulfil the training needs of our community. Due to a variety of reasons some training opportunities are still only available at a specific time or place, or simply do not scale well enough to satisfy the demand.

Several projects, organisations and initiatives continue to put their efforts into developing and sharing of teaching HPC resources. The EuroLab-4-HPC project establishes training in form of online courses¹⁴. The Barcelona Supercomputing Centre (BSC) aims to develop a professional training curriculum [4]. Both XSEDE¹⁵ (The Extreme Science and Engineering Discovery Environment) and PRACE¹⁶ (Partnership for Advanced Computing in Europa), serving similar purpose in US and Europe, provide a variety of training opportunities and resources to make HPC more accessible to researchers across many scientific domains. The SciNet Certificate Program¹⁷ provides course material for three major Computer Science topics: Scientific Computing, High Performance Computing, and Data Science each with several subcategories. Whereas the initiatives such as the HPC University¹⁸, the Carpentries (the HPC Carpentry especially)¹⁹ and ACM SIGHPC Education Chapter²⁰ strive to make sharing of teaching resources and practices simpler, and more accessible to trainers and practitioners alike.

Finally, many companies offer training courses making their products and services easier to use by their end-users, e.g. NVIDIA offering courses in deep learning²¹ or Arm offering courses on various arm technologies²².

8 CONCLUSION

The HPC Certification Program allows the re-use of existing content but also makes it possible to create a new ecosystem in which HPC centres, research labs, academic institutions and commercial companies could offer the best of their teaching material. According

⁸<https://www.epcc.ed.ac.uk/msc>

⁹<https://www.liverpool.ac.uk/study/postgraduate-taught/taught/big-data-msc/overview/>

¹⁰<https://www.mhpc.it/>

¹¹<https://www.maths.tcd.ie/hpcmsc/>

¹²<https://masters.fib.upc.edu/masters/meri-high-performance-computing>

¹³<http://univ-cotedazur.fr/education/training?AIHPC18&lang=en>

¹⁴<https://www.eurolab4hpc.eu/>

¹⁵<https://portal.xsede.org/web/xup/training/overview>

¹⁶<http://www.training.prace-ri.eu/>

¹⁷<https://www.scinethpc.ca/scinet-certificate-program/>

¹⁸<http://hpcuniversity.org>

¹⁹<https://hpc-carpentry.github.io/>

²⁰<https://sighpceducation.acm.org/>

²¹<https://www.nvidia.com/de-de/deep-learning-ai/education/>

²²<https://www.arm.com/support/training>

to our proposal, the existing and newly created teaching resources should be marked accordingly to indicated which skills they cover. In the future, the program may provide means to register and reference existing content of third-parties allowing users to browse the skills and navigate to teaching material.

8.1 Benefits

The program brings multiple benefits to everyone involved in HPC teaching and training. It's obvious that making clear what skills are required or recommended for a competent HPC user would be helpful to both the HPC service providers and practitioners. Training providers could bundle together skills that are most beneficial for specific user roles and scientific domains, which would allow practitioners to browse through skills to quickly identify and learn the skills required to perform their tasks. The variety of training offered within our community makes finding the right resources more complicated than it should be – the certification program will provide useful information where the desired skills could be learned. The examination confirming that a certain set of competencies has been acquired makes the learning process more complete and meaningful.

By participating in the program the HPC training providers can increase the visibility of their teaching opportunities and share their resources more effectively. The mapping of the skills defined by the program onto the existing training materials should also help to identify any potential gaps and improve the integrity of offered training. Finally, the certificates recognised by the whole HPC community simplify inter-comparison of independently offered courses and provide additional incentive for participation. Overall, the flexibility of the program allows to construct more personalised and just-in-time pathways to learning about HPC.

ACKNOWLEDGMENTS

We are thankful for the contributions and discussions with the members of the HPCCF and for the contributions made by the PeCoH project particularly Nathanael Hübbe for developing tools. PeCoH was supported by the German Research Foundation (DFG) under grants LU 1353/12-1, OL 241/2-1, and RI 1068/7-1.

REFERENCES

- [1] Trevor Hussey and Patrick Smith. 2008. Learning outcomes: a conceptual analysis. *Teaching in higher education* 13, 1 (2008), 107–115.
- [2] Julian Kunkel, Kai Himstedt, Nathanael Hübbe, Hinnerk Stüben, Sandra Schröder, Michael Kuhn, Matthias Riebisch, Stephan Olbrich, Thomas Ludwig, Weronika Filinger, Jean-Thomas Acquaviva, Anja Gerbes, and Lev Lafayette. 2019. Towards an HPC Certification Program. *Journal of Computational Science Education* (01 2019), 88–89. <https://doi.org/10.22369/issn.2153-4136/10/1/14>
- [3] Neil C Rowe. 2004. Cheating in online student assessment: Beyond plagiarism. (2004).
- [4] Maria-Ribera Sancho. 2016. BSC best practices in professional training and teaching for the HPC ecosystem. *Journal of Computational Science* 14 (2016), 74–77.
- [5] DM Williamson. 2011. Good Practice Guide on Writing Aims and Learning Outcomes. *Queen Mary, University of London, Pub 7282* (2011).