D7.3

Third (final) report on the activity of the High-Level support services (third year)

Mariella Ippolito, Nicola Spallanzani, Francisco Ramirez, Marnik Bercx, and Giovanni Pizzi

Due date of deliverable: 30/09/2022
Actual submission date: 18/10/2022
Final version: 07/10/2022

Lead beneficiary: CINECA (participant number 8)
Dissemination level: PU - Public

Document information

Project acronym:	MaX
Project full title:	Materials Design at the Exascale
Research Action Project type:	European Centre of Excellence in materials modelling, simulations and design
EC Grant agreement no.:	824143
Project starting / end date:	01/12/2018 (month 1) / 30/09/2022 (month 46)
Website:	www.max-centre.eu
Deliverable No.:	D7.3

Authors: Mariella Ippolito (Cineca), Nicola Spallanzani (Cnr), Francisco Ramirez (EPFL), Marnik Bercx (EPFL), and Giovanni Pizzi (EPFL)

To be cited as: M. Ippolito et al., (2022): Third (final) report on the activity of the High-Level support services (third year). Deliverable D7.3 of the H2020 project MAX (final version as of 07/10/2022). EC grant agreement no: 824143, CINECA, Bologna, Italy.

D7.3 Third (final) report on the activity of the High-Level Support services (third year)

Executive Summary

The MAX outreach activities are largely based on offering services of interest to a broad range of users, both academic and industrial. These services have been continuously improved during the duration of the project, to always guarantee users maximum support for exploiting at the best MAX codes on different kinds of architectures.

The main services offered are: i) the code and domain specific support for MAX codes, ii) the high-level support, addressing more complex problem such for example the development of new features in MAX codes, iii) the support for the whole suite of AiiDA tools (i.e., aida lab, quantum mobile, materials cloud) iv) the production of container-based deployment strategies to allow for easy access to codes and workflows of the MAX ecosystem.

The domain specific support is offered in two different ways: on the one side, the support is offered by each code community through forum/mailing lists and some other channel (e.g., github threads, slack), where the person in charge is always a specialist acting on behalf of MAX; on the other side, the support is explicitly operated through the MAX Help-desk. The latter offers both advanced support as well as a consultation for the utilisation of codes. Support to industries is also provided within this framework. As regarding the high-level support requests, they can be sent directly to the code communities support of each MAX code, via forums/mailing lists, filling a dedicated form on the website, or to the MAX Help-desk. Moreover, both the Help-desk and each code community also offer support for using the containerised versions of MAX codes.

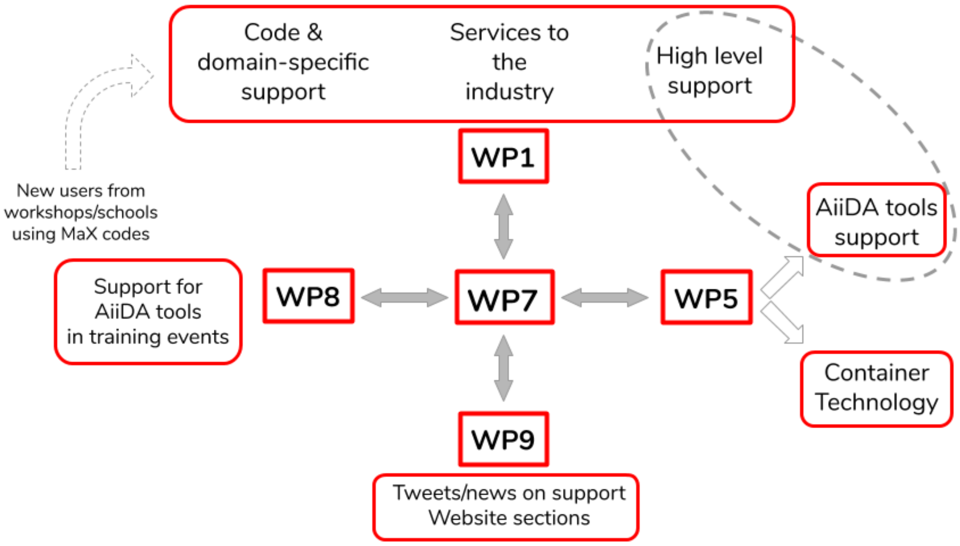

Fig. 1 Network of collaboration between MAX WPs.

The organisation of MAX services has been increasingly improved during the course of the project, in particular defining a well-organised network of collaborations among WPs in order to make all the support services more efficient and to give them more visibility. This network is represented in Fig. 1. The code and domain-specific support and the high-level support activities are provided by each MAX code community in synergy with the WP1. The Aiida tool support and the deployment of container-based technologies are provided mainly within the frame of WP5. The training events organised by WP8 represent a good opportunity to advertise the MAX support activities and for the engagement of new users. Moreover the support activities are mainly advertised also by tweets and news in the MAX website in the frame of WP9.

As regarding the support operated by each MAX code community in the period M24-46, we addressed via forum or mailing list about 1500 threads, corresponding to a total of about 3000 emails/post, in addition we managed more than 1300 gitlab threads and 64 face-to-face support activities. Looking at the entire duration of the project, considering also the support offered via gitlab and the face-to-face support, we solved more than 7000 queries.

In Fig. 2 it is reported the percentage of threads (arrived via forum or mailing list) for each MAX code with respect to the total.

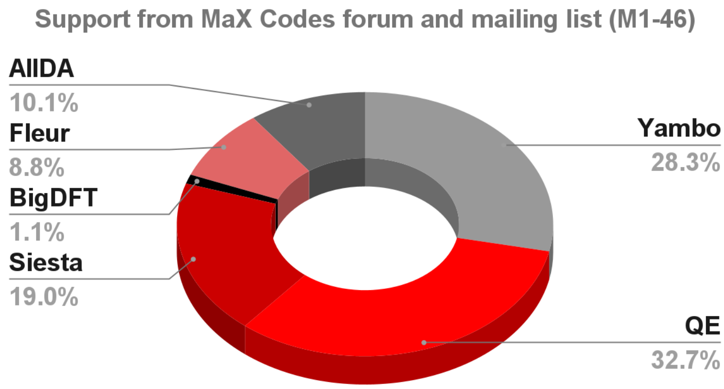

Fig. 2 Percentage of email/forum threads for each MAX code M1-46.

Many additional requests arrived also via slack, which proved to be the favourite channel for support requests for AiiDA and BigDFT users: more than 84000 slack messages were exchanged in the period M24-46, and a total of more than 140000 slack messages in the entire duration of the project M1-46. Slack is largely used inside MAX for interactions within a WP or among WPs.

In addition to the support operated by each code community, an activity is operated by the MAX Help-desk, that also embraces some non-MAX codes largely used in the material science community, such as lammps, vasp, and cpmd. In the covered period (M24-46) 377 service support actions of the Help-desk and 1092 related emails are reported. Instead, in the entire span of the project 840 support actions and 2794 emails have been performed. Of all these, about 51% were for MAX codes, and the remaining 49% for non-MAX codes.

In Fig. 3 the percentage of threads managed by the MAX Help-desk for each MAX code for the period M1-46 is reported . Most of the support requests were for the Quantum Espresso (QE) suite, but a lot of requests also for CP2K, SIESTA, and Yambo were submitted, too.

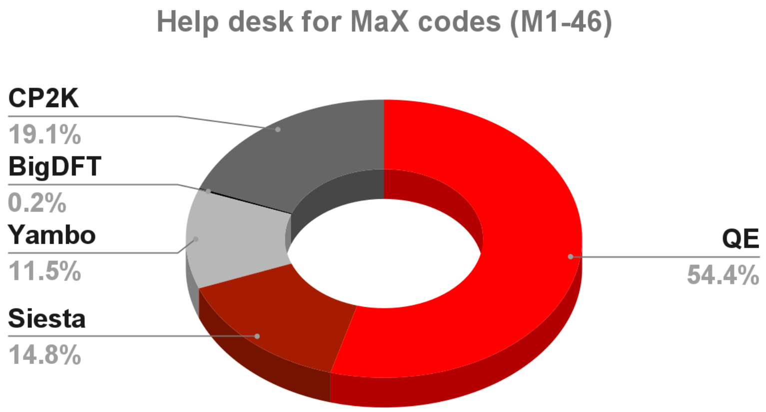

Fig. 3 Percentage of threads managed by MAX help desk referring to MAX codes for the period M1-46.

As for the high-level support, we addressed 60 support requests in the period M24-46, most of them regarding code usage on different architectures and requests for developing new features in MAX code, but there were also many requests for porting and benchmarking of MAX codes (particularly on GPUs). Looking at the entire duration of the project we addressed a total of 119 high-level support requests. The statistics of the high-level support for the period M1-46 are reported in Fig. 4.

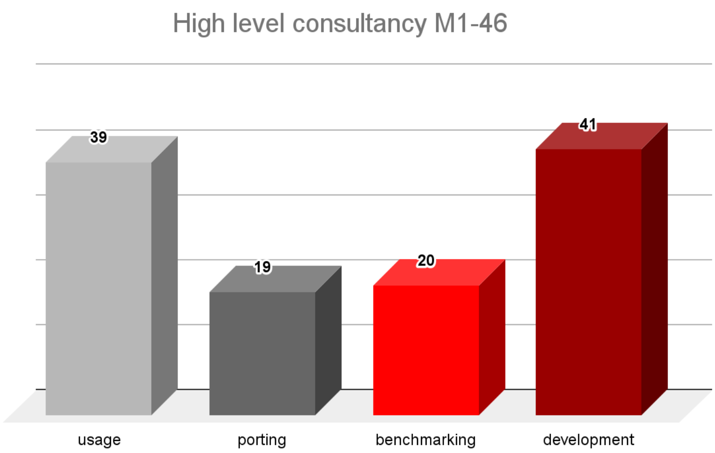

Fig. 4 High level consultancy request for MAX codes (M1-46).

For all the MAX codes we also provided one or more containerised versions, often benchmarked in the MAX supercomputer centres; the images for these are available in public repositories.

2. Code and domain-specific support

As already mentioned, code and domain specific support is operated on one side by the code communities, *via* forums/mailing lists, gitlab threads and slack messages with specialists acting on behalf of MAX, and on the other side *via* the MAX Help-desk service. The support activities carried out by the communities of the individual codes and the ones operated by the MAX Help-desk will be reported here separately.

Users can send a support request to MAX Help-desk by writing to the email address *support@max-centre.eu*. The messages are automatically forwarded to the Help-desk service operated through the CINECA's Trouble Ticket System (TTS) where, for each request, a ticket is opened in a queue dedicated to MAX. Also, the requests for codes belonging to the MAX area which are sent directly to the CINECA User Support team are classified in the MAX queue and taken in charge by the MAX staff. Depending on the level of desired support, the requests are addressed by MAX staff or, when necessary, taken over by other pertaining support services. Then they are tracked until the appropriate solution is given.

The MAX website has a dedicated Help-desk section[1] which details how to ask for MAX support and what kind of services are guaranteed. MAX support is advertised in the MAX website and in the various events organised by WP8.

2.1 Support activities

Typical domain-specific support actions can be distinguished in:

- investigation of the possible causes of a job failure and solution of the related problems. Typically this requires to re-build the code using different optimization parameters, or linking different libraries. In easy cases, it is only necessary to correct some input parameters or specifications of the computational resources in the batch script;
- evaluation of the MAX code performance on different architectures. For example, this is necessary to select the best architecture for a PRACE access proposal;

[1] http://www.max-centre.eu/services/max-help-desk

- analysis of a MAX code that behaves differently from documentation (e.g., non-converging algorithms);
- debugging of problems due to a specific code implementation (e.g., GPU, MPI, OpenMP versions);
- selection of the best code parameters that minimize the time to achieve a converged solution (e.g., input parameters or run time parameters);
- optimization of the computational setup to maximize the performance of a certain simulation on a given architecture (e.g., to find the best combination of number of nodes, number of cores, number of mpi processes/threads, etc.)
- support in the usage of different releases of MAX codes (user guidance about new *vs* deprecated features).

2.2 Report on the first level support activities and KPI breakdown

For this WP, the Key Performance Indicator (KPI) is given by the total number of support requests addressed by the CoE (trouble tickets, consultancy). In this regard, we collect the data related to the support activities operated by each MAX code community (Table 1) and the ones related to the MAX Help-desk activities (Table 2).

First, we focus on the support provided by each MAX code community. In Table 1, covering the period from M24 to M46, we present the different ways in which support was requested/provided since the support activity is organised in different ways for different MAX codes, e.g., forum, mailing list, gitlab issues, Slack messages. The "incoming threads" column presents the number of different requests arrived *via* email or in a forum. However, the total number of threads is not representative of the support activity carried out because each thread is typically followed by other emails/posts until the problem is solved. In order to take this into account, we report in the third column the total number of emails/posts for each MAX code.

Beside the official channels, we provide support in several other ways.

- "Face-to-face" support refers to assistance provided in a more personal and direct way. Rather than relying on email, the "face-to-face" service is given by a personal interaction, e.g. *via* telephone, *via* skype call or being physically present in the same location.
- "Gitlab threads" are the issues opened on the gitlab space of a given MAX code. Here we consider only the requests managed by MAX personnel. To use gitlab it is necessary to be registered and the subscription is free.

- "Slack messages" column contains the total number of messages shared in the Slack channel. Slack is a collaboration hub where conversations take place through channels organised by topic, project, team, or any issue considered relevant. In order to use Slack, it is necessary to be registered and in this case the registration is free. As it can be seen from the table, both AiiDA and BigDFT use Slack intensively to support their users. In the last months, the QE community started to use Slack for supporting users, too. Moreover Slack is largely used inside MAX also to interact within a WP or among different WPs.

MAX Codes	Number of incoming threads	Total number of emails/posts	Gitlab threads	Number of face-to-face support activities	Number of slack messages
QE	434	610	102	8	1367
SIESTA	212	573	89	12	-
Yambo	504	1054	-	26	-
BigDFT	-	-	17	6	40400
FLEUR	153	627	243	1	-
AiiDA	110	410	851	11	53797
Total	1413	3274	1302	64	84314

Table 1. Report M24-46 of the support operated by the single code communities and by MAX personnel.

In order to highlight the great support offered to MAX users during the entire duration of the project we collected in Table 2 also the whole support actions performed by each MAX codes community for the period M1-46.

MAX Codes	Number of incoming threads	Total number of emails/posts	Gitlab threads	Number of face-to-face support activities	Number of slack messages
QE	982	1401	163	8	1367
SIESTA	570	1465	141	12	-
Yambo	850	1847	-	39	-

BigDFT	34	48	85	20	71659
FLEUR	263	1021	470	3	0
AiiDA	304	1141	3182	79	75297
Total	3003	6923	4041	161	148323

Table 2. Report M1-46 of the support operated by the single code communities and by MAX personnel.

In Table 3, we report the support activity operated by the MAX Help-desk in the period M24-M46 (this service also includes the support given for non-MAX codes largely used in the material science area). We report explicitly the support activity for Vasp and LAMMPS, that are the most used softwares, while we gather in "Others" label the support given for all the other codes (e.g., WANNIER90, ADF, CPMD).

The first column refers to the number of different threads and the second to the total number of received emails for each code. As already mentioned, each thread corresponds to a ticket opened in the MAX queue at the CINECA's TTS.

Codes	Number of Incoming threads	Total number of emails
QE	112	316
SIESTA	25	64
Yambo	25	75
CP2K	42	133
Vasp	83	267
LAMMPS	53	141
Others	37	92
Total	377	1092

Table 3. Report M24-46 of the support operated by MAX Help-desk. The label "Others" refers to non-MAX codes in addition to Vasp and Lammps (i.e ADF, CPMD, WANNIER90).

In the period M24-46 covered by this data collection, 189 service support actions were performed by the Help-desk for MAX flagship codes and 563 related emails. Additionally, there have been 173 service support actions directed to other codes and 502 related emails, for a total of 377 actions and 1092 emails. There were no support requests for AiiDA, BigDFT, and FLEUR since users of these three codes refer mainly to support provided by the codes community.

In Table 4 we also report the support actions performed by MAX Help-desk for the whole duration of the project.

About 2800 total actions were performed by MAX Help-desk.

Codes	No. of Incoming threads	Total no. of emails
QE	231	757
SIESTA	63	213
Yambo	49	164
BigDFT	1	2
CP2K	81	271
Vasp	218	791
LAMMPS	106	341
Others	91	255
Total	840	2794

Table 4. Report M1-46 of the support operated by MAX Help-desk.

2.3 FAQ section

In order to reduce the costs connected to answer repeatedly to the same kind of requests for the MAX support team, we prepared a frequently asked questions (FAQs)[2] section on the MAX website. The section does not regard MAX in general, as we have chosen to constantly enhance and update the FAQs on the individual MAX codes. This activity is guaranteed by the leaders of each code. The links to the FAQ page of the various MAX codes are listed in the FAQ section, created in collaboration with WP9. This helps the users find immediate answers to frequent problems of relevance to several MAX codes. Moreover, in order to offer answers

[2] http://www.max-centre.eu/services/faq

related to more practical problems such as accounting problems, different environments and scheduler on different architectures in the Computing Centres involved in MAX, we also link the FAQ pages of these computing centres.

3. High level consultancy in materials science

In addition to the code and domain specific support MAX also provides a high level support targeting more complex problems related to MAX codes typically taking a long time to be solved and requiring to develop *ad-hoc* solutions for the customer, involving code development or a refactoring.

Support actions can be distinguished in four fields:

1. *code porting* on different architectures;
2. *code benchmarking* for evaluating MAX codes performances on different architectures. For example, this is necessary for selecting the best architecture to be requested in a PRACE access proposal or in other proposal;
3. *code usage,* regarding for example
   - debugging of problems due to a specific code implementation (e.g., GPU, MPI, OpenMP versions);
   - investigation of the possible causes of a job failure and solution of the related problems. In the same case, this can require to re-build the code using different optimization parameters, or linking different libraries. In most complicated cases the job failure can be caused by a bug and can be highly effort-requiring;
4. *code development*, comprising for example:
   - development of new features in the codes;
   - support in adding packages to Quantum Mobile and for AiiDA lab off site deployment.

3.1 Organisation of the high level support

The requests for high level support can be sent directly to the code communities of each MAX code, *via* forums/mailing lists, or to the MAX Help-desk (*support@max-centre.eu*). Within the MAX website, it is also possible to send specialistic support requests directly through the "High-level consultancy" subsection in the "Services" section by filling a dedicated form. This procedure is simple and immediate to submit and requires no registration. In this case the requests arrive at the MAX Help-desk where the request is handled directly, if the Help-desk is able to respond. This is the case, for example, of requests for porting and benchmarking of MAX codes, optimization of calculations on certain architectures, help in setting up some QM or QM/MM calculations with MAX codes and codes in the material science area in general. In some cases, on the other hand, it is necessary to interact with the support community of the code. Many requests require the implementation of new features in the codes or some bug fixing. When so, the request is scaled up to WP1 where a referent of the code takes charge of the request. However the request remains formally in charge of the Help-desk that serves as contact point for both the user and WP1 expert till the solution.

In most cases the high-level support requests are sent directly to the MAX code communities. The support activity is organized in different channels in the different codes, e.g., forum, mailing list, GitLab threads or Slack messages, thus the support requests may arrive in many different ways. To keep track of all these requests, we use a shared document in the MAX G-drive, and used by WP7 and WP1 staff. The details of the request are all reported: name and institution of the applicant, code, person in charge of the request, date of the request, state of the request (open/closed), and a brief description. This document is continuously updated and monitored by the MAX Help-desk.

3.2 Report for the High Level Support

The high-level support activity mainly regards the development of new features and new tools in MAX codes, and besides the code usage and delivering solutions from porting and benchmarking. In the table below, the support activities addressed in the period M24-46, both provided by the code communities and/or operated by the MAX Help-desk, are collected. For each request, we report the owner of the request, code, applicant, date of the request, status (i.e., closed or ongoing), and a brief description of the activity. When the "applicant" field is populated with "several users", it means that similar requests concerning the entire code community were posed by several users.

The received requests are too many to individually comment, thus here we will comment only the most relevant. A more detailed description of support activities related to the whole suite of AiiDA tools will be reported in the next section.

The great part of the high-level support requests addressed by code communities was on the development of new features and new tools and on bug fixing. Thanks to the problems pointed out by users we were able to identify many bugs in QE GPU, which required a lot of effort in solving.

A lot of work was devoted by MAX Help-desk to the porting and benchmarking of some MAX codes, and some other materials science codes, on GPUs. Particular effort was dedicated to the porting and benchmarking of CP and PH packages of the QE suite on Marconi100 (the Power9 CINECA architecture equipped with Nvidia V100 GPUs), being of interest for several users. Further benchmarks were required on Marconi100 for the PWscf package, too, in order to obtain an extreme scaling (up to 800 GPUs). A lot of effort was also required for CP2K porting and benchmarking on both Marconi100 (M100) and Galileo100 (G100), a CINECA Intel Cascade Lake architecture. CP2K failed on G100 due to a memory issue, caused by a bug related to the Intel compiler from version 2018 onwards, and only with intel2022 the problem disappeared. This bug has been permanently fixed in the version 9.1.

In order to better serve the materials research community, MAX offers support for some additional codes that are largely used by material science communities, in particular Vasp and Lammps. These two codes are largely used also by the Eurofusion (EF) community, whose high-level support team often interacts with the MAX Help-desk. This happened in particular for the porting of Lammps and Vasp on different architectures, and for performing related benchmarks.

A particularly relevant activity is the support given to Air Liquide[3] (AL) industry. In the first year of the project AL asked MAX support for running high-pressure calculations with QE. To this extent, an agreement for continuous support was signed at the beginning of 2020 by both CINECA and CNR. CINECA supported them for QE calculations and QM/MM calculations using the code PWQMMM interfacing LAMMPS with QE up to the second year. For a further year, CNR provided dedicated training on the use of the AiiDA framework together with the QE suite of codes by way of customized tutorials to exploit scientific use cases of interest to AL. The main result was the development of an AiiDA plugin containing the AiiDA work-chain that collects automatically all the steps of the workflow and keeps track of the provenance of all the input and output data.

Owner of the request	Applicant	Code	Date of request	Status	Type of request
M. Ippolito (CINECA)	Luca Sementa (IPCF-CNR Pisa)	CP2K	09/2021	Closed	Porting and benchmarking on G100 and M100

[3] https://www.airliquide.com/industry

M. Ippolito (CINECA)	Luca Sementa (IPCF-CNR Pisa)	CP2K	09/2021	Closed	Memory issue due to intel compiler (from version 2018 to 2021)
M. Ippolito (CINECA)	Simona Achilli (Politecnico di Milano)	SIESTA MAX-rel	03/2021	Closed	Building SIESTA MAX-rel on M100 with ELSI and psml support. Support for input set up.
M. Ippolito (CINECA)	Danila Amoroso (CNR-SPIN)	vasp	02/2021	Open	Bug identification in vasp/6.2.0 gpu version for magnetic calculation
M. Ippolito (CINECA)	Duc Nguyen (UKAEA)	vasp	08/2022	Closed	Vasp/6.3.2 porting and benchmarking on Marconi SKL
M. Ippolito (CINECA)	Several users	qe	10/2021	Closed	Cp benchmarking on M100
M. Ippolito (CINECA)	Ivan Carmineo (SISSA)	qe	03/2022	Closed	Pw benchmarking on M100 up to 800gpus
M. Ippolito (CINECA)	Ivan Marri (UniMORE)	qe	05/2022	Closed	Pw benchmarking on G100 for IscraB proposal
L. Bellentani (CINECA)	Several users	qe	03/2022	Closed	PH porting and benchmarking on M100
M. Ippolito (CINECA)	Davide Tisi (Sissa)	lammps/DeepMD	06/2021	Closed	Building lammps with DeepMD toolkit on M100
Sebastiaan Huber (EPFL)	Brandon Wood (Lawrence Livermore National Lab, USA)	AiiDA	09/2020	Ongoing	Request to add SSL support for connecting to the RabbitMQ message broker

Sebastiaan Huber (EPFL)	Brandon Wood (Lawrence Livermore National Lab, USA)	AiiDA	09/2020	Ongoing	Request to help getting \`aiida-nwchem\` up and running on the clusters of LLNL.
Elsa Passaro (EPFL)	Raimon Fabregat (EPFL)	Materials Cloud	28/12/20	Closed	Support to add contributed Discover app to the Materials Cloud page
Marnik Bercx (EPFL)	Miki Bonacci (Modena University)	AiiDA	03/2021	Ongoing	Provide support for SG15 pseudopotentials in \`aiida-pseudo\` so they can be run with the \`aiida-quantumespresso\` protocols.
Marnik Bercx (EPFL)	Ine Arts (University of Antwerp)	AiiDA	04/2021	Ongoing	Help with the installation of AiiDA on a data server connected to their HPC clusters.
Emanuele Bosoni (ICMAB-CSIC) Sebastiaan Huber Giovanni Pizzi (EPFL)	Samuel Poncé, Austin Zadoks (EPFL)	AiiDA	02/2020	Ongoing	Support in the creation of the Abinit implementation of the AiiDA common workflows, a set of AiiDA workflows to compute various material properties with a common interface among different quantum engines.
Emanuele Bosoni (ICMAB-CSIC) Sebastiaan Huber Giovanni Pizzi (EPFL)	Pezhman Zarabadi-Poor (University of Bath)	AiiDA	02/2020	Ongoing	Support in the creation of the ORCA implementation of the AiiDA common workflows, a set of AiiDA workflows to compute various material properties with a common interface among different quantum engines.
Emanuele Bosoni (ICMAB-CSIC) Sebastiaan Huber Giovanni Pizzi	Conrad S Johnston (Queen's university	AiiDA	02/2020	Ongoing	Support in the creation of the NWChem implementation of the AiiDA common workflows,

(EPFL)	Belfast)				
Emanuele Bosoni (ICMAB-CSIC) Sebastiaan Huber Giovanni Pizzi (EPFL)	Kristjan Eimre (empa)	AiiDA	02/2020	Ongoing	Support in the creation of the Gaussian implementation of the AiiDA common workflows, a set of AiiDA workflows to compute various material properties with a common interface among different quantum engines.
Emanuele Bosoni (ICMAB-CSIC) Sebastiaan Huber Giovanni Pizzi (EPFL)	Bonan Zhu (University College London)	AiiDA	02/2020	Ongoing	Support in the creation of the CASTEP implementation of the AiiDA common workflows, a set of AiiDA workflows to compute various material properties with a common interface among different quantum engines.
Emanuele Bosoni (ICMAB-CSIC) Sebastiaan Huber Giovanni Pizzi (EPFL)	Espen Flage-Larsen (SINTEF Industry Oslo)	AiiDA	02/2020	Ongoing	Support in the creation of the VASP implementation of the AiiDA common workflows, a set of AiiDA workflows to compute various material properties with a common interface among different quantum engines.
Marnik Bercx (EPFL)	Eric Macke (University of Bremen)	AiiDA	05/2021	Ongoing	Provide protocols for existing Hubbard-U work chain and help with the integration of Hubbard-V work chain in the AiiDA Quantum ESPRESSO plugin.
Giovanni Pizzi, Aliaksandr Yakutovich (EPFL)	Nicola Colonna (PSI)	AiiDAlab	05/2021	Ongoing	Installation and configuration of AiiDAlab for PSI

Giovanni Pizzi, Aliaksandr Yakutovich, Simon Adorf (EPFL)	Sandvik Coromant	AiiDA, AiiDAlab	03/2021	Closed	Support on installation of AiiDA and AiiDAlab in industrial setup
Giovanni Pizzi, Chris Sewell, Emanuele Bosoni	Oleg Rubel (McMaster University, Canada)	AiiDA	07/2021	Ongoing	AiiDA-Wien2K plugin support (development, bug fixes). U
Giovanni Pizzi	Dave Wecker (Microsoft)	AiiDA, AiiDAlab	04/2021	Ongoing	Deployment of AiiDA and AiiDAlab on Azure. U
Giovanni Pizzi	Christophe Galland (EPFL)	Materials Cloud	05/2021	Ongoing	Support for adding a new discovery section to Materials Cloud. D
Marnik Bercx	Mihail Petrov (UA)	AiiDA	8/10/2021	Ongoing	Deployment of AiiDA on university infrastructure. U
Paolo Giannozzi (CNR-IOM), Pietro Delugas(SISSA)	Matteo Coccoccioni (UniPavia) Iurii Timrov, Nicola Marzari (EPFL)	Quantum ESPRESSO	04/2020	Work ongoing Last update 4/2021	Improvement of LDA+U input and incorporation of new EPFL developments in Q.E.
Pietro Delugas (SISSA)	Aldo Romer (West Virginia Un.)	QE	05/2020	Ongoing	Improving output compatibility of projwfc.x with PyPROCAR functionalities
I. Carnimeo, S. Baroni, P. Giannozzi, S. De Gironcoli (SISSA)	Marco Aldinucci, Jacopo Colonelli (UniTO)	QE	10/2020	Ongoing	Usage of QE in docker based workflows using StreamFlow
Paolo Giannozzi Stefano Baroni Ivan Carnimeo Pietro Delugas	Jacob Gavartin Alexandr Fonari (Schrodinger)	QE Maestro-Schroedinger	02/2021	Ongoing	General Support and interface between Quantum ESPRESSO and Maestro GUI

(SISSA)					
Ivan Carnimeo Pietro Delugas (SISSA)	Filippo Spiga Louis Stuber (NVIDIA)	Quantum ESPRESSO	01/2021	Ongoing	Test, profile and maintenance of NVIDIA docker for Quantum ESPRESSO in NGC
Paolo Giannozzi (SISSA)	Alessandro Fortunelli, Luca Sementa (IPCF-CNR Pisa)	Quantum ESPRESSO	04/2021	Ongoing	New method for Virtual-Crystal Approximation in Quantum ESPRESSO
Paolo Giannozzi (SISSA)	Marco Pala (CNRS Paris-Saclay)	Quantum ESPRESSO	03/2021	Ongoing	Interfacing with Quantum ESPRESSO and public release of a quantum transport code
Pietro Delugas (SISSA)	Simone Piccinin, Matteo Farnesi (CNR-IOM)	Quantum ESPRESSO	02/2021	Ongoing	Fixing issues in image parallelism for NEB calculations in Marconi 100
N. Spallanzani A.Ferretti (CNR-NANO)	N. Marzari G. Pizzi (EPFL)	QE. Wannier90	10/2020	Ongoing	Optimization of memory management introducing a Scalapack routine in pw2wannier90 post-processing tool of Quantum ESPRESSO
N. Spallanzani A. Ferretti(CNR-NANO) M. Bonacci (UniMoRe)	Federico Iori (Air Liquide)	AiiDA QE	03/2020	Closed 03/2021	Training concerning the use of the AiiDA package together with the Quantum ESPRESSO suite of codes. Building of scripts and AiiDA-WorkChains addressed to the specific needs of Air Liquide
N. Spallanzani A. Ferretti (CNR-NANO)	EoCoE	QE	03/2021	Ongoing	Optimization in pw2wannier90 post-processing tool of Quantum ESPRESSO
N. Spallanzani (CNR-NANO)	Several users	Yambo	04/2021	Closed 04/2021	Support for the usage of the Yambo's container during the CECAM "Virtual school on electronic excitations in solids

					and nanostructures using the Yambo code"
A. Ferretti (CNR-NANO)	A. Marrazzo, M. Peressi (UniTS)	unfold/ QE	04/2021	Closed 07/2021	Compilation and usage of unfold.x, including atomic projections, with latest versions of QE.
Daniel Wortmann, Gregor Michalicek (FZ-JUELICH)	Felix Dushimineza (FZ-JUELICH)	FLEUR	12/2021	Closed 02/2022	Extend FLEUR functionality for the inclusion in a TEM image simulation infrastructure and provide related FLEUR usage support
Gregor Michalicek (FZ-JUELICH)	Chun Law (FZ-JUELICH)	FLEUR	02/2021	Closed 04/2021	Provide scripts and extend documentation for the extraction and visualisation of (projected) band-structure and density-of-states data in the Fleur-generated banddos.hdf file.
Daniel Wortmann, Gregor Michalicek (FZ-JUELICH)	Henning Janssen (ETH)	FLEUR	10/2020	Closed 12/2020	Increase IO performance and reduce amount of IO for density-of-states calculations employing the tetrahedron method with large k-point sets.
Gregor Michalicek (FZ-JUELICH)	Dongwook Go, Alexander Neukirchen (FZ-JUELICH)	Fleur	02/2021	Closed 04/2021	Implement option to generate calculation setups with reduced symmetry detection to simplify postprocessing in the context of interfacing to several external tools.
Pablo Ordejón (ICN2)	Simona Achilli (Politecnico di Milano)	SIESTA	06/2021	Closed 09/2021	Parallelization over orbitals of the Unfold utility in SIESTA.
Alberto García (CSIC) Nick Papior (DTU)	Pol Febrer (ICN2)	SIESTA	04/2022	Closed 05/2022	Guidance on the use of the hybrid MPI/OpenMP environment.

Federico Pedron (ICN2) and Alberto Garcia (CSIC)	Ernane Freitas (ICN2)	SIESTA	05/2021	Ongoing	Enhancement of the QM/MM subsystem in SIESTA.
Alberto Garcia (CSIC)	Yann Pouillon (SIMUNE) and Vladimir Dikan (CSIC)	SIESTA (libgridxc)	05/2021	Closed	New additions to the libgridxc library to handle the semilocal part of hybrid functionals, and to compute derivatives with respect to the gradient of the density.
Luigi Genovese	Maxwell Gisborne (Plymouth U.)	BigDFT	03/2022	Closed	Deployment of BigDFT SDK container.
Luigi Genovese	Samuel Dechamps (U. Louvain)	BigDFT	04/2022	Closed	Compilation of BigDFT code from source.
Fulvio Paleari (CNR-Nano)	Lorenzo Sponza (ONERA-CNRS)	Yambopy	06/2022	Closed	New feature: Brillouin zone expansion from QE databases.
Fulvio Paleari (CNR-nano)	Simone Grillo (Roma-Tor Vergata U.)	Yambopy	06/2022	Closed	Support for new lattice geometries and scissor operator in exciton band plots.

Table 5. List of High-level support activities.

In Fig. 5 we report the percentage of request types in the period M24-46. The most common requests were for code usage and development of new features and tools in the MAX codes. Instead, the requests for porting and benchmarking decreased from the second year, when, due to the starting production of the Power9 Marconi100 architecture, there was a strong interest in porting and benchmarking the materials science codes on that new architecture.

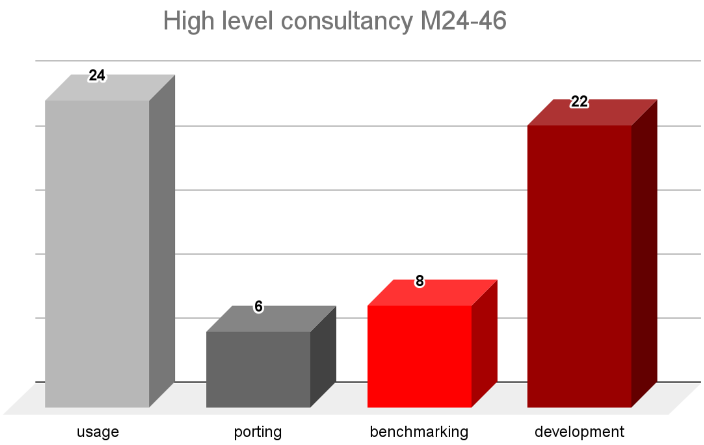

Fig. 5 High level consultancy request for MAX codes (M24-46).

3.3 Services to the industry

MAX also develops personalised consultancy to industries for specific and targeted needs, as outlined in the dedicated section in the MAX website[4]. Indeed, the very positive experience with Air Liquide helped us define a protocol for industry support: any request made by an industry is assessed by the MAX expert team and a plan is prepared. Examples of requests range from performing single accurate calculations to developing tailored workflows or intensive high-throughput simulations.

Some of the services provided are:

- support on optimal choice of parameters and parallel configuration to run user calculations for best performance (porting and/or benchmarking);
- ad-hoc solutions for users, possibly including custom code developments tailored at running MAX codes in non-standard conditions;
- dedicated consulting for the implementation of new features and/or new post-processing tools in MAX codes;
- training on demand: while MAX training events are usually available free of charge, industrial partners may request dedicated training, with custom topics and calendar (both on and off their premises).

[4] http://www.max-centre.eu/services-industry

Support is offered irrespectively of the type of computational resources used, from private clusters to national resources, as well PRACE resources.

4. Support for the AiiDA Ecosystem

MAX offers support for the whole suite of AiiDA tools, which can be broadly divided into: the AiiDA core package[5], the AiiDAlab interface[6], the Quantum Mobile virtual machine[7], and the Materials Cloud website[8]. The support provided typically involves technical assistance on how to use the tools, the addition of packages and/or features on request, help in both setting and maintaining off site deployments, etc. Below, a summary of the recent relevant developments involving each of these platforms is given.

The AiiDAlab interface and the Quantum Mobile virtual machine are important tools in the application and implementation of WP5 tasks T5.2 and T5.3 (Turn-key solutions and Quantum-as-a-Service), while the Materials Cloud is a key component for task WP5 T5.4 (Data stewardship).

4.1 AiiDA Core - Automated Interactive Infrastructure and Database for Computational Science

AiiDA core is the base package of AiiDA, which ships with all the necessary functionalities to manage the execution of calculation and workflows while automatically keeping track of their data provenance. It can easily be expanded through its plugin system,[9] so that users can acquire pre-configured packages and don't need to manually set all up.

The latest major version release of the code (v2.0) included several improvements for an enhanced user experience and/or for the setting up of a restructured baseline that allows to properly tackle some of the more important remaining limitations. The most notable examples of this include:

- The storage architecture of the backend has been completely refactored in order to address issues reported by users: high data duplication, heavy repositories that could

[5] https://aiida.readthedocs.io/projects/aiida-core/en/latest/
[6] https://aiidalab.readthedocs.io/en/latest/
[7] https://quantum-mobile.readthedocs.io/en/latest/
[8] https://www.materialscloud.org/
[9] https://aiidateam.github.io/aiida-registry/

not be compressed, reaching the limit of nodes before filling the full disk space, limitations to inspect the content of export files, etc.

- The node *namespace* has been restructured to simplify the findability and accessibility to key methods. It is now organised into several sub-namespaces (attributes, extras, repository, etc.) only revealed in the pertinent context.
- The plugin entry point system has been improved to avoid running the reentry scan command, facilitating the process of plugin development (and installation).

For more information and a full list of changes, the changelog is distributed with the code and is available online in the github repository[10].

4.2 AiiDA lab - A platform for accessible Materials Simulations

AiiDAlab is a cloud platform based on jupyter and jupyterhub, that includes AiiDA and the plugins for most of the MAX flagship codes already pre-installed and ready to be used. Moreover, thanks to its "appmode" interface, simple but very powerful GUIs can be easily developed, and used to drive quantum simulations directly from the browser. For on-premises installations, besides using the version included in Quantum Mobile (see below), the AiiDAlab platform can be deployed in an OpenStack installation, either within a company or in a data centre.

Since the main purpose of the AiiDAlab platform is to provide a user-friendly interface to access the features of AiiDA and other computational tools, its whole development procedure is heavily driven by user feedback.

The first type of granted support activity concerns the deployment of the AiiDAlab. We created instructions and templates to simplify the process of deploying AiiDAlab in a Kubernetes cluster in general, and in particular Azure cloud resources, making it as simple as that to set up an AiiDAlab platform. This can be used for AiiDA tutorials and makes AiiDAlab able to run resource-consuming calculations with Azure resources, without additional supercomputing resources. The *aiidalab-launch tool* enables users to launch an AiiDAlab instance directly on their local computer or workstation across all major operating systems, including native support for the ARM64 architecture.

The help provided last year in the deployment of AiiDAlab instance, on the computing clusters of public/private external institutions, included the deployment in Materials

[10] https://github.com/aiidateam/aiida-core/blob/v2.0.0/CHANGELOG.md

Modeling MarketPlace project hosted at DigitalOcean, and a representative of Sandvik Coromant via the public mailing list.

The computer and code setting procedures in the app are refactored and make the computational resource configuration simple. The aiida-code-registry is created to store the configuration templates of various computational centres and clusters, which can be directly loaded to set up these computational resources.

It is worth mentioning a new "Molecules" app developed at Empa: it allows to compute electronic properties of small molecules (such as ionisation potential, excitation energy, natural orbitals, and more) using Gaussian code. There is also a new SSSP pseudopotential toolkits app developed in EPFL, available on the AiiDAlab app store: it runs complex workflows to verify the quality of pseudopotential, and also provides tools to generate and select the pseudopotential for simulation.

4.3 Quantum Mobile - a VirtualBox machine

Quantum Mobile is a VirtualBox machine that comes with AiiDA and a set of commonly used quantum codes preinstalled. Its use does not require any initial setup, so it is easy to use for teaching purposes and to locally run calculations using AiiDA, or even to use AiiDA to manage remote computational resources with a consistent environment. The virtual machine already includes the majority of the MAX flagship codes and many more, together with the corresponding AiiDA plugins and workflows.

4.4 Materials Cloud - An online hub for accessible Materials Simulations

The Materials Cloud website acts as a hub of tools and data in computational materials science, shows redirects to the webpages for downloading AiiDA and Quantum Mobile or accessing the supported deployments of AiiDAlab. Further, it has the core purpose of hosting all data related to materials science research in the Archive section. Selected datasets can be explored and visualised in the Explore and Discover sections.

An important part of the support provided by the Materials Cloud team involved the adaptation of external contributions of curated research databases to be displayed with tailored visualisation tools in the Discover section of the website. The most recent additions are:

- Pyrene-based metal organic frameworks curated data set, containing atomic structures of experimentally synthesised pyrene-based metal-organic frameworks. The underlying AiiDA database is also available in the corresponding Explore section.
- Molecular vibration explorer, an interactive tool for the exploration of the vibrational spectra of 5k molecules and their applications in surface-enhanced spectroscopies (IR, Raman, SFG).
- Donor-acceptor copolymers for intramolecular singlet fission, an interface visualising the screening of 2944 conjugated donor-acceptor dimers for intramolecular singlet fission behaviour.
- New contributions to the Materials Cloud two-dimensional materials (MC2D) and 2D topological insulators databases.

Another important part of the support activities involves incorporating new computational tools into the Work/Tools section. In the last year, two new contributed tools were added as a service:

- A tool that calculates and visualizes the band structure of twisted mixed multilayer graphene in the k • p model.
- A machine-learning accelerated identification of exfoliable two-dimensional materials tool, which allows users to upload the bulk crystal structure in several standard formats (or to choose one from a set), and then layered structures are identified on geometrical criteria. After generating feature vectors representing the crystal structure, the tool uses a machine learning model to check if it can be exfoliated or has high binding energy.

5. Container technologies on HPC systems

The goal of this task is to produce container-based deployment strategies to allow for easy access to codes and workflows of the MAX ecosystem. This should allow computer centres to provide a uniform and well-defined environment that fully enables reproducibility, in a way that remains versatile enough to not impair the user interface and to be compatible with the automation of workflows.

Currently, AiiDA (and AiiDAlab) come with easily portable and redeployable containers (also on scalable technologies like kubernetes). Images for these are available in public repositories and the related links are available in the services section of the MAX website.

Also, all the other MAX codes provide containerised versions, often benchmarked in the MAX supercomputer centres: the exact status for each code in this regard is reported in the first subsection below.

The most relevant specific event that happened related to this topic was the creation of a Materials Cloud@HPC panel[11] that gathered a board of High Performance Computing experts from the major European HPC centres. This panel discussed strategies to address the challenges of containerisation in a common way among centres, that could be early adopters of the technology. So far, only one formal meeting has been organised, focused on the implementation of native support for running containerised codes directly into the AiiDA infrastructure, as well as the possible support for this feature inside of the various associated centres. The results are reported in the second subsection 5.2 below concerning running containerized codes with AiiDA.

5.1 Containerized codes

Containerised versions have been developed for all the MAX flagship codes and the support for the Singularity container is guaranteed in most of the HPC centres partner of MAX.

In the following, we report a list of the container technologies available for each MAX flagship code. All the related links can be found in the container technology section of the MAX website.[12] The MAX repository on DockerHub, with some of the most recent developments, is also available[13].

i. **Quantum ESPRESSO**: the Quantum ESPRESSO team provides a nvidia-docker image “nvcr.io/hpc/quantum_espresso” publicly available in the NGC catalog. The software is available from version 6.6a1 to the latest version (to date, version 7.1) with an installation that enables the GPU acceleration on Nvidia cards.

ii. **SIESTA:** the SIESTA team is developing a stable image of the SIESTA code. Currently, an image of a development version of the code is publicly available on the gitlab container registry of the SIESTA project .

iii. **Yambo:** the Yambo team provides a Docker image publicly available in docker hub. The “maxcentre/yambo” image provides a parallel installation of the executable files from version 5.0.1 to the latest (up to now 5.1.0). The container was successfully used in the hands-on sessions of the schools to provide the

[11] https://www.materialscloud.org/team (Materials Cloud@HPC Board section)
[12] http://www.max-centre.eu/services/max-container-technology-hpc-system
[13] https://hub.docker.com/u/maxcentre

students with a working version of the code, avoiding the complexity of the source installation.

iv. **FLEUR:** the FLEUR team provides an image in which they explain the usage of their code. The benefit is to provide the user with a fully functional "virtual environment" to use FLEUR and try out different tutorials. The image ”judft/fleur”, publicly available in docker hub, can be managed with both Docker or Podman.

v. **CP2K:** the CP2K team provides two Docker images developed outside the MAX consortium. The image “cp2k/cp2k” is publicly available in docker hub, and the image “nvcr.io/hpc/cp2k” is publicly available in the NGC catalog.

vi. **BigDFT:** the BigDFT team provides two Docker images, both of them publicly available in docker hub. The “bigdft/sdk” image provides containers for environments in which the development version of the code can be compiled and executed. With the tag developers can choose the proper environment for the compilation. The “bigdft/runtime” image is for BigDFT users who don't need to modify the code, as it is a smaller compiled version that should run on most platforms. It also includes jupyter notebook support with BigDFT bindings, to ease interactions with BigDFT.

vii. **AiiDA:** the AiiDA team provides three Docker containers, all publicly available in docker hub. The “aiida-core” image is available from version 1.1.0 to the latest version (to date 2.0.3). This image contains a fully pre-configured AiiDA environment particularly useful and handy for learning and testing purposes. A second image, “aiida-prerequisites”, contains all prerequisites allowing to run AiiDA. This repository adds PostgreSQL and RabbitMQ servers on top of the phusion base image. Additionally, it creates a system profile ready to set up AiiDA under it. The thread image “aiidalab/aiidalab-docker-stack” contains the Docker stack for deploying and starting the AiiDAlab (https://aiidalab.materialscloud.org). The limitation of using the AiiDA containers is that all data stored in a container will persist only over the lifetime of that particular container, unless you use volumes (in the AiiDA documentation there are useful instructions to overcome this limitation).

5.2 Running containerized codes with AiiDA

Whenever an AiidA-run calculation requires the execution of an external code (e.g., a binary on a remote HPC cluster), the information related to that code has to be pre-configured inside a node that the calculation takes as input. Since remote computers usually have different architecture, different authentication methods, different job management, etc., part of this information stored has to be specific for the computer in which the code is installed. Therefore, even if the codes run the same version of simulation software, it is still necessary to create a separate code node for every remote computer.

The last years have seen increasing adoption of containerization technologies, included in the HPC domain, where executables are no longer installed directly on the target machine but are compiled once into a portable image and then run in an encapsulated environment.

This encapsulation of the full run-time environment, as well as the availability of global container registries, constitute a major step forward in terms of reproducibility. Moreover, storing the identifier of the container in the AiiDA graph makes it possible to directly re-run existing workflows without accessing the computer where it was originally executed.

Meanwhile, the containerized codes are less susceptible to the specifics of the remote computer, which simplifies the setup of the code for each new computer and reduces the amount of metadata that needs to be recorded.

The feature of running the containerized code as a first-class citizen in AiiDA is close to being finished. A first implementation of it performing manual code installation inside the containers has already been finalised and was tested on both Singularity and Sarus platforms. Before releasing a production-ready version, we are still looking into using conda to simplify the management of the containerized software.

5.3 Specific challenges of GPU containerization

The containerization of a GPU-aware code has been performed with BigDFT code, and included in the NVidia NGC repository[14].

The nvidia-docker platform has been employed to deploy the code in the suitable environment. We have verified that this technique enables the user to employ different versions of the CUDA drivers, depending on the host platform architecture. This container enables an efficient usage of the GPU resources leaving the details of the compilation of the

[14] https://catalog.ngc.nvidia.com/orgs/hpc/containers/bigdft

executable to the container builders. Also, other container environments like singularity have been tested with the same approach.

6. Conclusions

The organization of support activities has given very positive feedback during the entire duration of the project. We improved year by year the organization and the quality of the MAX services, in particular by enhancing the high-level support, thanks to a strong collaboration between WP7 and WP1, and introducing a well-defined protocol for supporting industrial requests. This has been possible thanks to a well defined network of collaborations between WPs.

During the project, the MAX Help-desk managed about 840 support requests providing more than 2800 actions. More than 7000 queries were solved from MAX codes communities. Additional 120 requests were addressed by high-level support.

The outcome of this amount of support activities shows how important for the code development the MAX user feedback is: thanks to a continuous fruitful interaction, we were able to improve the codes, to identify and solve many bugs, and to address other related problems.

Concerning container technologies, all the MAX codes provided containerised versions, often benchmarked in the MAX supercomputer centres, whose images are now available in public repositories. We also created a MAX repository containing most of these MAX codes containers. Moreover Singularity has been installed in most supercomputer centres partner of MAX and the support for the usage of containerized version of MAX codes is also guaranteed.

ANNEX 1 - List of Acronyms

Acronym	Description
CoE	Centre of Excellence
CUDA	Compute Unified Device Architecture
EPFL	Ecole Polytechnique Fédérale de Lausanne
FAQs	Frequently asked questions
GPU	Graphics Processing Unit
GUI	Graphical User Interface
MPI	Message Passing Interface
OpenMP	Open Multiprocessing
QE	QuantumESPRESSO
QM	Quantum Mechanics
QM/MM	Quantum Mechanics/Molecular Mechanics
SSSP	Standard solid-state pseudopotentials
TTS	Trouble Ticket System
WP	Work Package