D7.2

Second report on the activity of the High-Level Support services (second year)

Mariella Ippolito, Francisco Ramirez, Giovanni Pizzi, and Elsa Passaro

Due date of deliverable: 30/11/2020
Actual submission date: 30/11/2020
Final version: 30/11/2020

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) / 30/11/2021 (month 36)
Website:	www.max-centre.eu
Deliverable No.:	D7.2

Authors: M. Ippolito, F. Ramirez, G. Pizzi, and E. Passaro

To be cited as: M. Ippolito et al., (2020): Second report on the activity of the High-Level Support services (second year). Deliverable D7.2 of the H2020 project MaX (final version as of 30/11/2020). EC grant agreement no: 824143, CINECA, Bologna, Italy.

Content

1. Executive Summary

The MAX outreach activities are largely based on offering services of interest to a broad range of users, both academic and industrial. We guarantee several services to support the use of MAX codes and of codes in the materials science area in general. Support is offered in two different ways: on one side, the support is offered by each code community through forum/mailing lists, where the responsible 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 an advanced support as well as a consultation for the utilisation of codes. Support to industries is also provided within this framework.

In addition, MAX support activities aim to produce container-based deployment strategies to allow for easy access to codes and workflows of the MAX ecosystem.

The organization of support activities had already given positive feedback during the first year of MAX; in the second year we tried to further improve this organization, in particular for the high-level support and the industrial support. Under this respect we have better defined the 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.

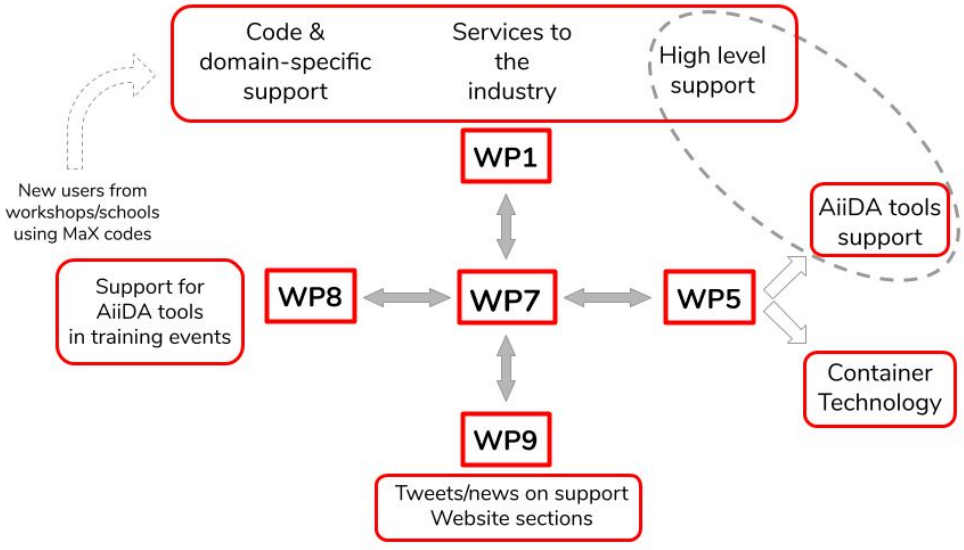

Fig. 1 Network of collaboration between MAX WPs

In the period M12-24 we addressed more than 1000 domain-specific support requests, considering both the threads managed by each code community and the ones managed by

the MAX Help-desk, for a total of 2828 emails. Many additional requests arrived *via* gitlab and *via* slack. As for the high-level support, we addressed about 30 support requests, most of them regarding porting and benchmarking of MAX codes. This is a consequence of the starting production of the new Cineca Tier-0 architecture, an IBM machine based on Power9 processors accelerated with NVIDIA GPUs named Marconi100.

2. Code and domain-specific support

Code and domain specific support is operated on one side by the code communities, *via* forums/mailing lists, with specialists acting on behalf of MAX, and on the other side *via* 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 directed to this address are forwarded automatically to the Help-desk service operated through the CINECA TTS system 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 solution of the problem.

The MAX website has a dedicated Help-desk section[1] which details how to ask for MAX support and what kind of service is guaranteed. MAX support is advertised not only on the MAX website but also in the various events organized by WP8: in fact many new MAX users are former attendees of MAX schools/workshops or Hackathons.

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 for selecting the best architecture to be requested in a PRACE access proposal;
- analysis of a MAX codes that behaves differently from documentation (e.g. non-converging algorithms);

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

- 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. 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

The Key Performance Indicator (KPI) for this WP 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).

In Table 1, covering the 2nd year of support activities from M12 to M24, we report several columns to distinguish the different ways in which support was requested/provided (the support activity is organized in different ways for different MAX codes, e.g. forum, mailing list, Slack messages). The "incoming threads" column represents 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.

Beyond the support offered through the official channels, there are support activities provided with several other methods.

- "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 organized 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 as a way of supporting their users. Slack is largely used inside MAX also to interact within WPs or between 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	276	405	31	-	-
Siesta	211	495	26	-	-
Yambo	185	399	-	4	-
BigDFT	-	-	23	3	16150
Fleur	64	227	167	-	-
AiiDA	110	408	1422	64	17168
Total	846	1934	1669	71	33318

Table 1. Report M12-24 of the support operated by the single code communities and by MAX personnel.

Comparing the values reported in Table 1 for the period M12-24 with data of the previous year (M1-12), it is apparent that the number of support requests and emails/posts handled by each code community has increased significantly.

In Table 2 we report the support activity operated by the MAX Help-desk in the period M12-M24 (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. ADF, CPMD, WANNIER90). The first column refers to the number of different threads, 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 CINECA system.

Codes	Number of Incoming threads	Total number of emails
QE	44	208
Siesta	11	57
Yambo	9	40
CP2K	19	82
Vasp	55	288
LAMMPS	28	124
Others	24	90

Total	194	889

Table 2. Report M12-24 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 M12-24 period covered by this data collection, 76 service support actions were performed by the Help-desk for MAX flagship codes and 377 related emails. Additionally, there have been 107 service support actions directed to other codes and 502 related emails, for a total of 183 actions and 879 emails. There were no support requests for AiiDA, BigDFT and Fleur: users of these three codes refer mainly to support provided by the codes community. Help-desk support requests were slightly lower than last year (210 support requests arrived in the period M1-12), however the number of emails has greatly increased (the total number of email in the period M1-12 were 559). This means that much more work was needed to resolve these requests.

3. High level consultancy in materials science

MAX high level support targets 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:

- code porting on different architectures;
- evaluations of the MAX codes performances on different architectures. For example, this is necessary for selecting the best architecture to be requested in a PRACE access proposal;
- 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 a lot of effort can be necessary to solve it;
- development of new features in the codes;
- support in adding packages to Quantum Mobile and for AiiDA lab off site deployment.

3.1 Organization of the high level support

The requests for high level support can be sent directly to the code communities support 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, if the Help-desk has the necessary skills to respond to the support request, this is handled directly by it. This is the case, for example, of requests for porting and benchmarking of MAX codes, optimization of calculations on certain architectures, help in the set up of some QM or QM/MM calculations with MAX codes and codes in the material science area in general. In some cases the support request can not be handled by the MAX Help-desk alone but it is necessary to interact with the support community of the code. In fact some support requests may require to make changes to a MAX code, this may be the case in which the users ask to implement new features in the code or to fix some bugs. In these cases 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 manages contacts with the user and interacts continuously with the WP1 until the request is closed.
In most cases the high level support requests are sent directly to the MAX code communities. The support activity is organized in a different way for each MAX code, 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 the requests that arrive to the different MAX code communities, there is a document in the MAX G-drive shared by WP7 and WP1 where all the requests for specialistic support are entered. All the details of the request are reported: the name and institution of the applicant, the code, the person in charge of the request, the date of the request, the state of the request (open/closed) and a brief description of the request. 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 code usage, delivering solutions from porting and benchmarking and development of new tools in MAX codes. In the table below we report all the support activities, the ones provided by the communities of the individual codes and the ones operated by the MAX Help-desk.

The high level support requests for the period M12-24 are reported in Table 3. For each request we reported the owner of the request, the code, the applicant, the date of the request, the status (i.e. closed or ongoing) and a brief description of the activity. For some requests in the "applicant" field we reported "several users": in fact, many users came up with similar requests concerning the entire code community.

The received requests are too many to comment on each of them, thus here we will comment only on the most relevant ones. Support activities related to the whole suite of AiiDA tools instead will be reported in the next section, as it requires a more detailed discussion.

By the end of April 2020, the new Cineca Tier-0 architecture started its production phase: it is an IBM machine based on Power9 processors accelerated with NVIDIA Vora V100 GPUs and it is named Marconi100. A lot of work was devoted by MAX Help-desk to the porting and optimization of all the MAX codes, and materials science codes in general, on this new Power9 architecture. Also some MAX code communities (e.g. QE, Yambo, SIESTA) worked on the porting and benchmarking of their own codes on Marconi100. In this regard, valuable help was also provided by the NVIDIA staff. In order to facilitate communication between the MAX Help-desk and the NVIDIA staff, and to make all the faced problems traceable, we created a project on GitLab dedicated to the porting of material science code to this Power9 architecture.

A lot of work was done to find the best configuration to compile QE-GPU on Power9 and the best setup to maximize performance. This required a strong collaboration between MAX Help-desk (WP7) and the code developers (WP1). In this regard, Slack was used as the main communication channel and frequent online meetings were also held. Thanks to the problems reported by users we identified many bugs in QE-GPU and a lot of effort was needed to solve them.

A large effort was devoted in performing CP2K benchmarks on different architectures with GPUs, i.e. Intel broadwell (bdw) nodes equipped with NVIDIA K80 GPU and Power9 processors accelerated with NVIDIA V100 (Marconi100). Strong scaling and weak scaling curves were obtained for different case studies. These benchmarks were performed to support the user for an IscraB proposal, the project has been submitted. Many benchmarks were performed also with cp2k+SIRIUS.

In order to better serve the materials research community, MAX offers support also to some additional codes that are largely used by material since communities, such as in particular Vasp and Lammps. These two codes are largely used also by the Eurofusion (EF) community and often the MAX Help-desk interacts with the EF high level support. In the last year this happened in particular for the porting of Lammps and Vasp on M100, and for performing related benchmarks.

Regarding Vasp, a lot of effort was devoted also on fixing bugs in Vasp/6.1.1 related to magnetic calculations with constraints. These bugs were spotted from version 6.1.0 onwards and emerged thanks to the report of some MAX users whose calculations on Power9 produced incorrect energies in output and in some cases hang. Once the bugs were identified, these were communicated to the code developers and to the Nvidia staff who took care of the porting of the code on GPUs with openACC. These bugs were fixed and the new version 6.1.2 of Vasp was released.

Another relevant activity is the support given to an industry, AirLiquide[2], for running high-pressure calculation with QE. After a positive feedback for the support provided during the first year of the project an agreement for continuous support was made on 31-01-2020

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

with both CINECA and CNR. We keep supporting them for QE calculations and QM/MM calculations using the code PWQMMM code interfacing LAMMPS with QE.

The experience of the support activity provided to AirLiquide - and the consequent industrial agreement for a continuous support - has traced the path for a well-defined protocol in supporting industrial requests, as explained in the section "Services to the Industry".

Owner of the request	Applicant	Code	Date of request	Status	Type of request
M.Ippolito (CINECA)	Ivan Girotto (ICTP)	cp2k	12/2019	Closed	CP2K benchmarks on x86 GPU architectures for an Iscra B Proposal at CINECA
M. Ippolito (CINECA)	I. Girotto (ICTP)	cp2k	06/2020	Closed	Porting and benchmarking on Marconi100
M. Ippolito (CINECA)	I. Girotto (ICTP)	Cp2k SIRIUS	09/2020	Ongoing	Building cp2k+SIRIUS on Marconi100 and benchmarking it
M. Ippolito (CINECA)	S. Meloni (University of Ferrara)	lammps	03/2020	Closed	LAMMPS benchmarks on GPU architectures for Prace Proposal
M. Ippolito (CINECA)	Serhiy Mochalskyy (Max-Planck-Institut)	lammps	05/2020	Closed	Porting and benchmarking on Marconi100
M. Ippolito (CINECA)	Matteo Busato (University of Udine)	cpmd	05/2020	Closed	Porting and benchmarking on x86 and Marconi100 for Iscra B proposal
M. Ippolito (CINECA)	Duc Nguyen-Manh (University of Oxford)	Vasp	05/2020	Closed	Porting and benchmarking on Marconi100
M. Ippolito (CINECA)	Giacomo Giorgi (University of Perugia)	Vasp	07/2020	Closed	Support for solving Vasp problems on Marconi100 related to different versions of PGI compiler

M. Ippolito (CINECA)	Federico Iori (Air Liquide)	Qe, PWQMMM	03/2020	Closed	Support for high pressure QE calculations and qm/mm calculations with PWQMMM
M. Ippolito (CINECA)	Several users	Qe-GPU, siesta, yambo	05/2020	Closed	Porting and benchmarking on Marconi100
P. Bonfa (CNR) P.Delugas (SISSA) F.Affinito (CINECA)	Several users	Qe-GPU	05/2020	Closed	Porting and benchmarking on Marconi100
P.Delugas (SISSA) P. Bonfa (CNR)	Several users	Qe, Qe-GPU	05/2020	Closed	Bug Fixing
Paolo Giannozzi (CNR-IOM), Pietro Delugas (SISSA)	Matteo Coccoccioni (UniPavia) Iurii Timrov, Nicola Marzari (EPFL)	Quantum Espresso	04/2020	Ongoing	Improvement of LDA+U input and incorporation of new EPFL developments in QE
Oscar Baseggio, Stefano Baroni (SISSA)	N. Vast, O. Motorny (LSI, Paris) Iurii Timrov (EPFL), Andrea dal Corso (SISSA)	Quantum Espresso	02/2020	Closed	EELS spectroscopy with the Sternheimer approach
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
N. Spallanzani A. Ferretti (CNR)	M. Celino F. Buonocore (ENEA, EoCoE)	QE, wannier90	12/2019	Ongoing	Optimization of the pw2wannier90 post-processing tool of Quantum ESPRESSO
A. Marini (CNR)	Valerie Smejkal (TU-Wien)	Yambo	02/2020	Ongoing	Usage of the real-time module of Yambo

A. Ferretti D. Varsano N. Spallanzani (CNR)	N. Marzari, M. Borelli (EPFL)	Yambo	05/2020	Closed	Provision and generation of scaling and performance data of Yambo on M100 for PRACE application
A. Ferretti D. Varsano (CNR)	M. Palummo (UniRoma2)	Yambo	05/2020	Closed	Porting of the Yambo code on Marconi100
A. Garcia (CSIC)	Several users	SIESTA	05/2020	Closed	Porting and benchmarking on Marconi100
Giovanni Pizzi (EPFL)	Pablo Piaggi (Princeton)	Materials Cloud	09/03/2020	Closed	Support to add contributed tools to the Materials Cloud page.
Giovanni Pizzi (EPFL)	Kristiāns Čerņevičs, Valeria Granata , Oleg V. Yazyev (EPFL)	Materials Cloud	01/05/2020	Closed	Support to add contributed tools to the Materials Cloud page.
Sebastiaan Huber (EPFL)	Brendan Wood (Lawrence Livermore National Lab, USA)	AiiDA	05/2020	Ongoing	Implement feature to allow configuration of all message broker settings
Sebastiaan Huber (EPFL)	Brendan 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)	Brendan Wood (Lawrence Livermore National Lab, USA)	AiiDA	09/2020	Ongoing	Request to help get AiiDA up and running on the clusters of LLNL. As a first proof-of-principle we are trying to run Quantum ESPRESSO calculations through the `aiida-quantum espresso` plugin.
Luigi Genovese	Michal Krompiec (Merck)	BigDFT	06/2020	Closed	Interface and API of PyBigDFT in HPC clusters

Table 3. Report M12-24 for the high level support activities

In Fig.2 we report the percentage of request types in the first and second MAX year. During the first year the largest number of requests were for development of new features and tools in the MAX codes. In the second year the largest number of requests were instead for porting and benchmarking: this is mainly due to the starting production of the new Power9 architecture Marconi100 and related requests for porting and benchmarking material science codes.

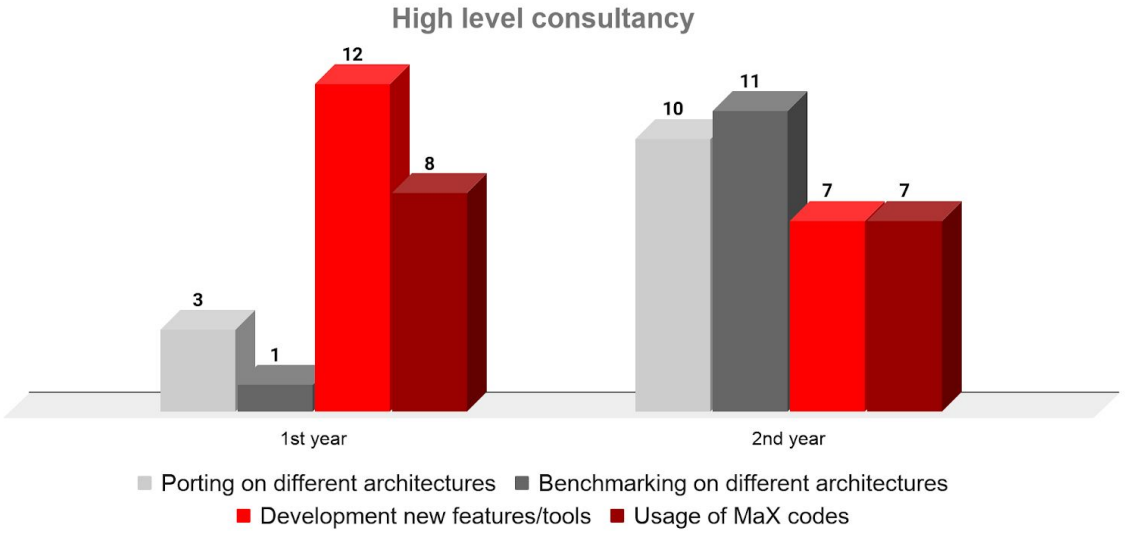

Fig. 2 High level support activities in the first and second MAX year, percentage for type of requests.

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[3]. Indeed, the very positive experience with AirLiquide helped us define a protocol for industry support. If any request is made by an industry, the MAX expert team evaluates it and proposes the best path to follow. Examples range from performing single accurate calculations to developing tailored workflows or intensive high-throughput simulations.

Same of the services provided are:

- support on optimal choice of parameters and parallel configuration to run user's calculations for best performance (porting and/or benchmarking);
- ad-hoc solutions for users, possibly comprising custom code developments tailored at running MAX codes in non-standard conditions;

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

- 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).

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

4. Support for the AiiDA Ecosystem

MAX also offers support for the whole suite of AiiDA tools, which can be broadly divided into: the AiiDA core package[4], the Materials Cloud website[5], Quantum Mobile virtual machine[6], and the AiiDAlab interface[7]. This support typically involves technical assistance on how to use the tools, the addition of packages and features on request, help in setting and maintaining off site deployments, etc. Below there is a summary of the recent relevant developments involving each of these platforms.

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

AiiDA core is the base distribution 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[8] so that users can acquire pre-configured packages and don't need to manually setup all. The most recent version of the software (version 1.x) has been documented and published in the "Scientific Data" journal of Nature[9].

The support activities related to AiiDA core are mostly organized through two channels: the aiida mailing list[10] and the github repository[11]. The mailing list is mainly a tool for questions related to AiiDA usage, both for the core package and any of its plugins. The github

[4] http://www.aiida.net
[5] https://www.materialscloud.org/
[6] https://www.materialscloud.org/work/quantum-mobile
[7] https://www.materialscloud.org/work/aiidalab

[8] https://aiidateam.github.io/aiida-registry/
[9] Huber, S.P., Zoupanos, S., Uhrin, M. et al. AiiDA 1.0, a scalable computational infrastructure for automated reproducible workflows and data provenance. Sci Data 7, 300 (2020). https://doi.org/10.1038/s41597-020-00638-4
[10] http://www.aiida.net/mailing-list/
[11] https://github.com/aiidateam/aiida-core

repository, besides being the place where the code is hosted, also has the "issues" sections where all development related requests are posted (bug reports, feature requests, etc).

Other relevant support activities might arise from direct contact with any of the current developers of the code. Last year's most relevant was a collaboration with Lawrence Livermore National Lab from which a multi-aspect support was requested. Their main objective was to deploy AiiDA to run on their clusters, along with further development-related activities of implementing more levels of configurability and adding SSL support for the message broker, in order to address the usage of supercomputer centres with 2-factor authentication (the idea was to have a secure installation inside the computer centre network).

4.2 Materials Cloud - An online hub for for accessible Materials Simulations

The Materials Cloud website acts as the main centralized hub from where being redirected to the pages to download AiiDA and Quantum Mobile, or access the supported deployments of AiiDAlab. It also has the more specific purpose of hosting all data related to material science research (in its Archive section), as well as offering different visualization and exploration tools to easily navigate through that information (using the Explore and Discover sections). We recently (September 2020) published a paper[12] in which we present the platform and its main features. This activity is related to WP5-T5.4.

An important part of the support provided by the Materials Cloud team involves the adaptation of external contributions of curated research databases to be displayed with tailored visualization tools in the Discover section of the website. The most recent additions (the last three of which obtained with the AiiDA workflows, so the full provenance can be browsed in the Explore section) are:

- Kernel principal covariates regression, which illustrates a new method to calculate a low-dimensional representation that reproduces a large fraction of the intrinsic structural variability of data sets of structures.
- Automated high-throughput wannierisation, showing the validation results of an automated protocol for generating maximally-localized Wannier functions in a high-throughput framework.
- Covalent organic frameworks for carbon capture, with the interactive representation of 70k two- and three-dimensional covalent organic frameworks, screened for carbon capture.

[12] Talirz, L., Kumbhar, S., Passaro, E. et al. Materials Cloud, a platform for open computational science. Sci Data 7, 299 (2020) [DOI: 10.1038/s41597-020-00637-5]

- Metal-organic frameworks for carbon capture, with the interactive representation of 324,426 hypothetical Metal-Organic Frameworks (MOFs) screened for wet flue gas CO2 capture.

Another important part of the support activities involves helping to incorporate new externally designed tools to the Work/Tools section. In the last year, two new contributed computational tools were added to be offered as a service. The first is an environment finder, a tool for finding and analyzing atomic environments in crystal structures. The second is a Tight-Binding Electronic Transport Application (TBETA) for graphene nanoribbon junctions, to build and analyze electronic transport properties of those materials.

Finally, support activities might also involve making adaptations for compatibility with external protocols and interfaces. In this regard, it is worth noting that all structures in the Explore section are now also served *via* the Optimade API (v1.0 of which was just recently released[13]), a common API to query materials-science databases in a common language (Materials Cloud is one of the over 10 different worldwide databases involved in the effort).

4.3 AiiDAlab - A cloud 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 then used to drive quantum simulations directly from the browser. For on-premises installations, besides using the version included in Quantum Mobile, the AiiDAlab platform can be deployed in an OpenStack installation, either within a company or in a data centre. We are already implementing a viable model (including technical, support and possibly legal aspects) to deploy these services for internal company use, to encourage industrial uptake. This activity is related to WP5-T5.2 and T5.3.

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. Therefore, most of the latest updates and features, such as the revised QE app function or the revised home app landing page, introduce many modifications directly requested by users of the platform. A notable one is the addition of an OPTIMADE structure query widget to the default structure selection widgets, which increases the interoperability with this API (OPTIMADE already described in the Materials Cloud section).

[13] https://www.optimade.org/

A second type of support activities related to this tool, consists of the help provided in the deployment of AiiDAlab instances on the computing clusters of public or private external institutions. In the last year, the most relevant cases were the deployments in the Paratera clusters in China[14], and on the kubernetes cluster hosted by CESNET[15] in Czech Republic, provided *via* EGI/EOSC Hub[16].

It is also worth mentioning here the existing collaboration with the Empa research institute[17] . Indeed, an important amount of support has been provided to simplify and improve the AiiDAlab apps developed by them, either directly or by updating and adapting the underlying aiidalab-widgets-base. Sometimes, this also involved adapting other AiiDA plugin packages, such as aiida-quantumespresso and aiida-cp2k which are required for many of Empa's use cases. Finally, a 1-week workshop was offered specifically for the migration of the Empa apps after the release of AiiDA version 1.0.

4.4 Quantum Mobile— a VirtualBox machine that comes with AiiDA and a set of commonly used quantum codes preinstalled

The use of Quantum Mobile does not require any initial setup, therefore it can be used 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 machine already includes the majority of the MAX flagship codes and many more, together with the corresponding AiiDA plugins and workflows. You can find all instructions to download and set up the most recent version on the newly released website[18]. Developments of this are related to WP5-T5.2 and T5.3

Many of the support activities for Quantum Mobile involve preparing up to date releases to be used for special events, such as hands-on sessions for tutorials or workshops (in collaboration with WP8). In the last year, different versions of the virtual machine were used for the following:

- A “2-day AiiDA tutorial” that took place in Tokyo on December 2019[19]
- A “Computational School on Electronic Excitations” at the ICTP on January 2020[20]
- The “Wannier90 School” held online on March 2020[21]

[14] https://www.paratera.com/
[15] https://www.cesnet.cz/?lang=en
[16] https://www.egi.eu/
[17] https://www.empa.ch/
[18] https://quantum-mobile.readthedocs.io
[19] https://atztogo.github.io/AiiDA-tutorial-ISSP/
[20] http://indico.ictp.it/event/9018/
[21] http://www.wannier.org/events/school-2020-virtual-edition/

- The first "AiiDA virtual tutorial" in July 2020[22]

For the last of these events, the AiiDA virtual tutorial, a newly introduced "Cloud Edition'' of the virtual machine was used, in which pre-built images were deployed on either AWS and GCP servers so that these can easily be accessed and used remotely. The adaptations to this new modality proved to be specially convenient given the online nature of the event.

Further details on these activities will be reported in WP8.

5. Users engagement

Some of the MAX codes have a very large user-base, established well before the first MAX CoE. This is for example the case of QE, CP2K and Siesta. For activities related to these codes, the impact on the communities is immediate. It is important for the MAX Consortium to communicate such novel activities and put the user in contact with expert developers, in order to improve the results thanks to their feedback. For example only thanks to the active interaction with some users it was possible to identify and solve some bugs due to the GPU implementation of some MAX codes.

In the case of the most widely used MAX codes, users also take an active role in the support activities. In fact looking at the forum/mailing list of such codes for example we estimate that for QE about 20% of the answers are provided by non-developer users and for Siesta they are at least 10%.

A different perspective is provided by the activities of MAX in relation with younger/less used codes, that do not have to face the burden of legacy code, stratified along the years, and that have the possibility to modify external API more easily. This is for example the case of BigDFT. These activities may be associated with more intrusive modifications and/or experimental solutions, which contribute to strengthen the know-how of MAX consortium. Such outcomes, by their nature, would require more time to be widely diffused to the users' communities, and the expected timeline to gather feedback from them may go beyond the project's duration. Nonetheless, we believe that such activities would contribute to enlarge the basis of the know-how of the future generation of code developers in our field.

A separate discussion applies to AiiDA: although it is a recent code, it already has a very large user-base. Looking at the AiiDA plugin registry[23] it can be seen that many of the plugins are developed outside of MAX, with active developers that actually also provide user support -

[22] http://www.aiida.net/aiida-virtual-tutorial-july-2020/
[23] https://aiidateam.github.io/aiida-registry/

but most importantly, AiiDA teams are engaging the community in the support part, and even more in the development part, obtaining if possible stronger results.

5.1 Impact of the support activities on the broader MAX code communities

MAX support requests, especially the high level ones, often lead to results of interest to the entire community using the MAX codes: this is for example the case of the implementation of new features in a code or the resolution of a bug. In this case, the modifications to a MAX code are reported on the website of the code where for each new version typically are reported the list of the new features and the bugs fixed.

In order to reduce the costs connected to answer repeatedly to the same kind of requests for the MAX support team, a section dedicated to the most frequently asked questions (FAQs)[24] has been organized on the MAX website. Instead of creating a specific FAQs section for MAX, we have chosen to constantly enhance and update the FAQs of the individual MAX codes. This activity is guaranteed by leaders of each code. In addition, in collaboration with WP9 we have created a FAQs section in the services section of the MAX web site where we have reported the links to the FAQs page of the various MAX codes. This helps the users in finding immediate answers to frequent problems of relevance to several MAX codes. Moreover, in order to offer answers 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 list these computing centres' FAQ pages.

6. 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.
At the moment the support for the Singularity container is guaranteed in most of the HPC centres partner of MAX but we are evaluating also the newest SARUS as a better choice. Due to the lack of a single standard containerisation technology in computer centres (multiple technologies exist, based on Docker, like Shifter, Singularity, of SARUS) we have also investigated how to facilitate the creation of images that can be simply generated for various technologies.
Currently, AiiDA (and AiiDA lab) come with easily portable and redeployable containers (also on scalable technologies like kubernetes). Also, many of the MAX codes provide

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

containerised versions, often benchmarked in the MAX supercomputer centres, notably including Quantum ESPRESSO, CP2K, and BigDFT. Container images for AiiDA and BigDFT are available in public repositories and the related links are available in the services section of the MAX website. The singularity container for QE is available at Cineca but not yet available in a public repository.

Great advancements in this task are connected to three major events that happened last year. First, a meeting on the Sarus containerisation service (December 3, 2019, CSCS). This is a new generation system developed by CSCS that works as an alternative to Docker, Shifter and Singularity, addressing the issues that each of them has in a HPC context. During this meeting, the technology of Sarus was presented and it was discussed how to apply it to AiiDA deployments through both AiiDAlab and Quantum Mobile.
The second event was the "AiiDA Hackathon"[25] (February 17-18, 2020, CINECA). Although the main focus of the event was the development of plugins for AiiDA, Luigi Genovese (WP3, CEA) presented the results and conclusions on the previous year's experience in containerizing BigDFT. A focus group followed to discuss the experiences with AiiDAlab and Fleur and included topics such as usage in the context of testing and existing institutional practices.
Third, the "Access Abstraction to HPC Resources"[26] (October 22, 2020, virtual conference). In this event, several containerization strategies capable of accessing HPC resources were presented, both from the point of view of the cluster maintainers, and of the workflow management developers. Specific talks on AiiDA and AiiDAlab were given.

Finally, we are currently working to create a Materials Cloud@HPC panel that will gather a board of High Performance Computing experts from the major European HPC centres, with the objective of discussing strategies to address the challenges of containerisation in a common way among centres, in order to be early adopters of the technology. The resulting prescriptions will have to be compatible with the infrastructures of all major supercomputing centres (thus enabling the implementation of common shared protocols) and will focus on addressing any eventual complications that might prevent or complicate their application when reaching the computational speeds of the exascale.
Once this step is completed, we will be able to implement native support for running containerised codes directly into the AiiDA infrastructure. Different methods are currently being tested, and possible solutions have already been verified to work. Once the final consensus on the one acceptable for HPC computer clusters, we will finally be able to provide users with the ultimate solution to the goal of scientific reproducibility of research.

[25] https://events.prace-ri.eu/event/957/overview
[26] https://indico.psi.ch/event/8220/

7. Conclusions

The organization of the support activities has been improved with respect to the first year of the MAX project, and the network of collaborations between WPs has been improved.
The support activities have been given very positive feedback, considering both the support provided by the communities of each MAX code and the one provided by the MAX Help-desk.
During the period M12-24 we addressed more support activities with respect to the former period: this shows that, despite the Covid19, MAX support activities continued very efficiently, even if carried out mainly in smart working mode.
The outcome of many support activities shows how important the MAX users' feedback is: thanks to an active interaction with the users it has been possible to improve the codes, to identify and solve many bugs, and to address other related problems.
As for containers, Singularity has been installed in most supercomputer centres partner of MAX, though we are also evaluating the newest SARUS as an alternative. Many of the MAX codes provide containerised versions, often benchmarked in the MAX supercomputer centres, notably including Quantum ESPRESSO, CP2K, and BigDFT. The next step will be to prepare and make available the container images for all the other MAX flagship codes.