



H2020-FETHPC-01-2016



DEEP-EST

DEEP Extreme Scale Technologies

Grant Agreement Number: 754304

D7.1

Communication plan, toolkit and owned channels

Approved

Version: 2.0

Author(s): S. Eisenreich (BADW-LRZ)

Contributor(s): E. Suarez (JUELICH), A. Kreuzer (JUELICH), S. Höfler-Thierfeldt (JUELICH), A. Jackson (UEDIN), M. Ott (BADW-LRZ)

Date: 26.09.2018

Project and Deliverable Information Sheet

DEEP-EST Project	Project Ref. №: 754304	
	Project Title: DEEP Extreme Scale Technologies	
	Project Web Site: http://www.deep-projects.eu	
	Deliverable ID: D7.1	
	Deliverable Nature: Other	
	Deliverable Level: PU *	Contractual Date of Delivery: 31 / 10 / 2017
		Actual Date of Delivery: 31 / 10 / 2017
EC Project Officer: Juan Pelegrín		

* - The dissemination levels are indicated as follows: PU = Public, fully open, e.g. web; CO = Confidential, restricted under conditions set out in Model Grant Agreement; CI = Classified, information as referred to in Commission Decision 2001/844/EC.

Document Control Sheet

Document	Title: Communication plan, toolkit and owned channels	
	ID: D7.1	
	Version: 2.0	Status: Approved
	Available at: http://www.deep-projects.eu	
	Software Tool: Microsoft Word	
	File(s): DEEP-EST_D7.1_Communication_plan_v2.0	
Authorship	Written by:	S. Eisenreich (BADW-LRZ)
	Contributors:	E. Suarez (JUELICH), A. Kreuzer (JUELICH), S. Höfler-Thierfeldt (JUELICH), A. Jackson (UEDIN), M. Ott (BADW-LRZ)
	Reviewed by:	P. Petkov (NCSA), I.Schmitz (ParTec)
	Approved by:	BoP/PMT

Document Status Sheet

Version	Date	Status	Comments
1.0	31/October/2017	Final version	EC submission
2.0	26/September/2018	Approved	EC approved

Document Keywords

Keywords:	DEEP-EST, HPC, Exascale, communication, outreach, website, social media, communication toolkit
------------------	--

Copyright notice:

© 2017-2020 DEEP-EST Consortium Partners. All rights reserved. This document is a project document of the DEEP-EST project. All contents are reserved by default and may not be disclosed to third parties without the written consent of the DEEP-EST partners, except as mandated by the European Commission contract 754304 for reviewing and dissemination purposes.

All trademarks and other rights on third party products mentioned in this document are acknowledged as own by the respective holders.

Table of Contents

Project and Deliverable Information Sheet.....	1
Document Control Sheet	1
Document Status Sheet	2
Document Keywords.....	3
Table of Contents	4
List of Figures.....	5
List of Tables	5
Executive Summary	6
1 Introduction	7
2 Communication and Dissemination Plan	7
2.1 Strategy	7
2.2 Communication Toolkit.....	11
2.3 Owned Channels (digital)	14
2.4 News & Media Relations	16
2.5 Liaison with influencers & European collaboration	16
2.6 Dissemination of results: Scientific Publications, Deliverables & Technical Documentation	17
2.7 Conferences, events & trade fairs	18
3 Early Access Programme	19
4 Training & Education	20
4.1 Training topics.....	20
4.2 Training Schedule	21
4.3 Course material	22
5 Outlook	22
Annex A.....	23
A.1 Communication Toolkit for Download	23
List of Acronyms and Abbreviations	24

List of Figures

Figure 1: The DEEP-EST logo.....11

Figure 2: The DEEP projects logo11

Figure 3: DEEP projects flyer12

Figure 4: DEEP-EST flyer.....12

Figure 5: Newsletter template.....13

Figure 6: Presentation template - Title slide.....14

Figure 7: Presentation template – Content slide14

List of Tables

Table 1: KPIs..... 9

Executive Summary

The deliverable D7.1 Communication plan, toolkit and owned channels describes in detail:

- The DEEP-EST communication & dissemination plan, including the communication strategy, target groups, key messages and KPIs.
- The communication toolkit – which consists of the DEEP-EST flyer, the DEEP projects flyer, a template for presentations, a newsletter template, graphics & visuals as well as marketing collateral.
- The project's owned channels, which are online and/or updated by end of M4: the project website and the social media channels (Twitter, Facebook, LinkedIn).

Please note, this deliverable will detail the communication related activities of WP7, which are mainly carried out in Tk7.1. Still, the deliverable will include details on the communication and outreach related aspects of e.g. the open access strategy for scientific publications (Tk7.2), the early access programme (Tk7.3) or the training and education part (Tk7.4). Innovation and Exploitation, however, will not be discussed in detail in this deliverable – as they are strictly speaking not communication. A milestone is foreseen in the Description of Action (DoA) for this part of WP7 and is due at M18.

1 Introduction

Communication and dissemination of the results achieved in DEEP-EST is one of the project's main objectives. Measures will be taken to maximise the project's impact in the European and global HPC community and beyond. The basis for all the actions will be a comprehensive and concerted communication and dissemination campaign.

It is very important to understand that even though only three consortium partners share the main responsibilities in this work package, the planned activities will be executed in a concerted action by all partners. The three main responsible partners are BADW-LRZ for Tk7.1 communication, UEDIN for Tk7.3 the early access programme and JUELICH for Tk7.4 the training and education programme and contributions to the execution of the open access strategy as part of Tk7.2

This document will give a detailed overview on the communication strategy, the toolkit that has been developed and the status of the owned channels in section 2. An overview of the early access programme with a special focus on the communication-related aspects of it is given in section 3. Section 4 details the training and education programme – again with a focus on communication aspects as material on internal trainings will be available to an external audience and training and education events will be organised for external target groups as well. Last but not least, section 5 will give an outlook on next steps.

2 Communication and Dissemination Plan

Outreach activities by the DEEP-EST project will be grounded in a solid communication strategy. This strategy is broader in scope than the mere dissemination activities that focus on publishing the project results in scientific publications and at scientific conferences, in deliverables and technical documentation. We will proactively communicate on the project objectives, approach and outcomes to raise awareness and establish strong relations with key stakeholders (especially with potential users in academia and industry via a dedicated early access programme). In this way, collaboration with major players in the European HPC arena will be facilitated, thus contributing to the implementation of the SRA and finally to help position Europe as a technological leader in the race towards Exascale and HPC/HPDA convergence.

2.1 Strategy

DEEP-EST will continue to follow the communication strategy developed in DEEP/-ER, executed via a mix of owned, earned and paid channels supplemented with face-to-face encounters at conferences, events, and workshops.

- The project's **owned channels** form the hub of all communication activities. The focus here will be on the project's digital channels and include the project website (www.deep-projects.eu), the social media channels ([Twitter](#), [LinkedIn](#), [Facebook](#)) as well as the regular e-newsletters. Additionally, the project will have print material like the project flyers and potentially a case study booklet and a final brochure.
- **Earned channels** are (traditional) media outlets as well as channels of social media influencers. Earning publications in the form of e.g. interviews, news and background articles or likes, retweets and shares respectively, will be based on strong relations to journalists and social media influencers. These stakeholders will help increasing the

visibility and impact of the DEEP-EST project by mentioning, referring to, and reporting about the project to their large audiences.

- **Paid channels** can be e.g. advertorials, sponsored posts in social media or subscriptions to news platforms like e.g. Eurekalert and will be considered when deemed necessary and the cost-benefit ratio seems favourable for the project.
- **Events & Conferences** form the fourth pillar in our communication mix. Next to media transported outreach, face to face encounters with our target audiences will be key.

2.1.1 Objectives

The main objective of the communication efforts is to bring the overall Modular Supercomputing Architecture (MSA) concept, including its unique selling points, the project results, and lessons learned to the awareness of the defined target audiences. In concrete, WP7 has the following objectives:

- Communicate key features, objectives and outcomes of DEEP-EST to a wide range of stakeholders.
- Liaise with HPC industry, lobbying and special interest groups and standardisation bodies.
- Promote European partnerships for HPC development and attract industrial and academic HPC users to test the DEEP-EST prototype.
- Implement a targeted education and training programme for project internal staff as well as distribute the know-how gained in DEEP-EST to external communities.

2.1.2 Target Groups & Key Messages

The communication efforts of the DEEP-EST project will target mainly the following groups:

- a) HPC community (HW and SW developers, special interest groups like OpenHPC, EEHPC working group, young and female HPC researchers and students).
- b) Data centres as potential operators of future MSA-systems.
- c) Domain scientists and industrial HPC customers as potential users.
- d) Standardisation bodies (e.g. the OpenMP forum, the OpenMPI forum, the Big Data Value Association) and lobby groups (e.g. ETP4HPC).
- e) Multipliers like journalists and social media influencers.
- f) Political decision makers as well as the general public.

The key messages delivered to these target audiences will be in-line with the DEEP projects messaging when talking about general concepts, which are common to all three (DEEP, DEEP-ER, DEEP-EST). Naturally, these will be complemented with specific DEEP-EST messages. They will form the backbone of all our outreach activities be it flyers, brochures, slogans for trade fair booths, social media, interviews etc.

High-level messages that apply mostly for group (f) and partially (e):

- The DEEP projects develop a **new breed of innovative HPC systems**.
- The DEEP projects feature **high-impact HPC & HPDA applications** developed in Europe from important scientific and engineering fields and help make them fit for the future.
- The DEEP projects **foster the use and development of European HPC technologies**.

Our main target groups are technically-oriented stakeholders. Hence, it is most important to have clear messages for these groups (a-d).

- The **Modular Supercomputer Architecture (MSA)** is a blueprint for **heterogeneous HPC systems** supporting the divergent computation and data processing requirements

of **high performance computing** and **data analytics** with **highest efficiency** and **scalability**.

- Key arguments for data centre operators:
 - Ideal fit for centres running heterogeneous application mixes.
 - Fulfils the requirements of very diverse applications.
 - Significant reductions in time-to-solution and energy usage.
 - Improvements in system throughput for diverse workloads.
- **'True' Co-Design** that reaches from hardware to middleware/system software to tools to applications
 - Key talking point: Only via 'true' Co-Design it is possible to achieve Exascale performance on an application level.
- DEEP projects provide an **integrated programming environment**, using standards such as MPI, OpenMP and Tensorflow or similar frameworks for machine learning and data analytics.
 - Key talking point for application developers: **Ease of use** and portability.
 - Key talking point for data centre operators: **SLURM resource manager/scheduler** determine the optimal resource allocation for each workload and maximised the use of resources.

For group (e) the journalists and social media influencers the key messages will be tuned depending on how 'tech-savvy' they are. For the classic HPC outlets and social media influencer within the HPC community, the key messages for groups (a-d) apply better, for more general interest influencers, it is more about getting across the high-level messages and those related to the impact on society. On the contrary within group (f) the political decision-makers there might be candidates that are more into the details and hence would fall more into groups (a-d).

2.1.3 KPIs

The success of our communication campaign will be tracked with the following key performance indicators (KPI). We will monitor the success on a continuous basis and report on them in the progress and midterm reports. Based on how we are doing against the KPIs, we will adapt the communication campaign if needed.

Table 1: KPIs

Activity	Metric to measure impact	Goal	Timeline
Publications in peer-reviewed journals / scientific conferences	No of publications	15	Conferences: Potentially from Year 1 Journals: Year 2-3 and after the end of the project
Presentations at conferences	No of presentations	15	15 / year, whereas the first year might be slower than year 2 and 3
Project website	No of visits	45.000	Both figures: accumulated no /year
	Unique visitors	15.000	

Activity	Metric to measure impact	Goal	Timeline
Twitter	No of followers	ca. 1250	<u>Follower numbers</u> Above 800 in 1 st year Above 1000 in 2 nd year Above 1250 in 3 rd year
	No of impressions	ca. 400.000	<u>Impressions</u> Summed-up across whole project time frame
Facebook	No of likes for page	ca. 500	Ca 100 in year 1 Ca 250 in year 2 Up to 500 in year 3
	Type of audience	Mixed	<ul style="list-style-type: none"> • More female audience than e.g. on Twitter & LinkedIn • More diverse age group
LinkedIn	No of group members	150	
Press Releases & Newsletters	No of press releases issued	3	1 per year <ul style="list-style-type: none"> • Start of the project • End of the project • One in between, topic tbd (e.g. set-up prototype) Note: This does not include press releases by the partners on e.g. their products.
	No of newsletters issued	6	2 per year to be issued before e.g. big events like ISC and SC
	No of subscribers to newsletter	500	500 new contacts 300 contacts already exist ca 150 new contacts / year
Media Relations	No of interviews	8-10	ca 3 / year
	No of articles / news pieces in non-scientific publications	50	
Events	No of organised events	9-10	
Training & educational activities	No of internal trainings	3-5	1/year, more if needed
	No of external trainings & workshops	10	
	No of ext people reached	20 / event	

2.2 Communication Toolkit

The basis for all communication activities will be the communication toolkit that consists of flyers, a newsletter template, a template for presentations (PowerPoint and OpenOffice) as well as visuals (e.g. project map) and further marketing collateral. This toolkit will help us to ensure a uniform appearance of the project across all our channels (print and online).

2.2.1 Branding

DEEP-EST will be positioned as a further member of the “DEEP projects” family (following and extending DEEP and DEEP-ER). In general, the ‘DEEP projects’ branding will take precedence over the DEEP-EST one as the DEEP brand is already known to a wide audience. Still, logos have been designed for both DEEP-EST and DEEP projects following the established corporate design line.



Figure 1: The DEEP-EST logo



Figure 2: The DEEP projects logo

We will use the DEEP projects logo for the website, the social media channels, the e-newsletter as well as for collateral. Only the DEEP-EST project flyer will carry the DEEP-EST logo – as will e.g. more official documents like reports and deliverables. As is easily visible, the DEEP-EST logo transports the DEEP branding predominantly with the –EST claim below – meaning again, the DEEP brand takes precedence.

Additionally, we have put a focus on reworking the corporate design line towards a cleaner and clearer design. In the subchapters below it becomes visible that the corporate design now reaches across all materials print and online.

2.2.2 Flyer

Initially there will be two project flyers: One exclusively with content on the DEEP-EST project and one flyer on the DEEP projects in general to draw the evolutionary line. Eventually, there might be more flyers e.g. in the format of fact sheets on the prototypes or flyers / handouts informing on the early access program.



Figure 3: DEEP projects flyer



Figure 4: DEEP-EST flyer

2.2.3 Newsletter Template

We intend to send regular e-newsletters to the contacts in our database and aim for at least two issues per year. The newsletter will feature certain regular categories that will depend, however, on the availability of content. Categories include e.g. an editorial by project management, news stories on technical content, presentation of partner projects and researchers from the DEEP-EST project, event announcements. The newsletter template will look like this:



Figure 5: Newsletter template

2.2.4 Presentation Template

The presentation template will be used by all project colleagues for F2F meetings and as well as for reviews, conferences, workshops and events. It will be provided in 16:9 and in 4:3 in .pptx and in open office.

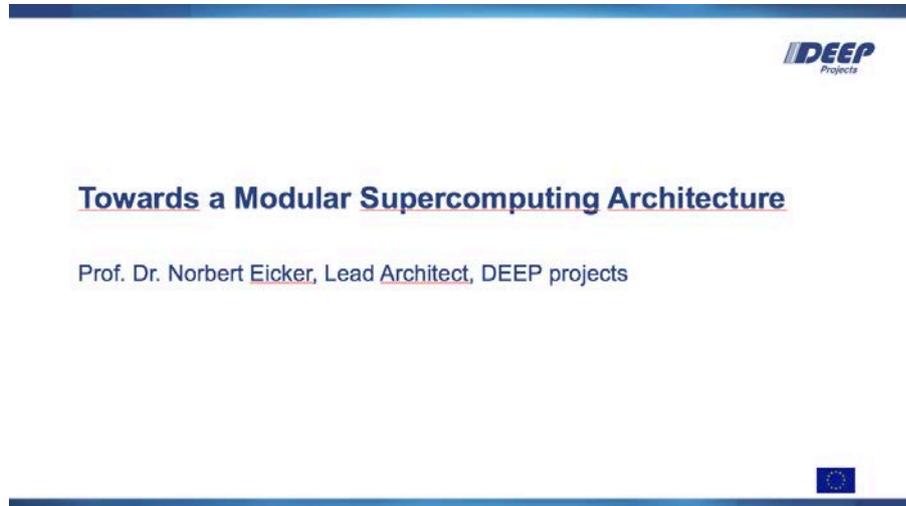


Figure 6: Presentation template - Title slide

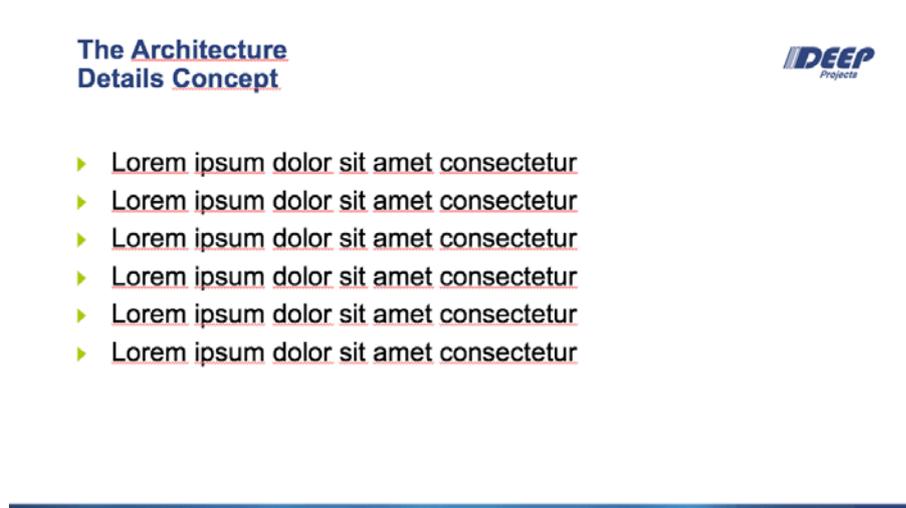


Figure 7: Presentation template – Content slide

2.2.5 Branding Materials

The communication kit will be complemented with buttons, stickers, T-shirts and potentially also give-aways. These will be developed within the next weeks.

2.3 Owned Channels (digital)

As mentioned in section 2.1, the outreach activities will focus on the owned channels and will include digital as well as print channels. This section deals with the online channels website and social media.

2.3.1 Website

The website will be the central information hub for DEEP-EST and the DEEP projects. This means, content will be updated and produced mainly for the website and then also pushed out via the social media channels and the e-newsletter.

In line with the branding (see section 2.2.1), the website will feature information about all DEEP projects. We have decided to continue and adapt the DEEP-ER project website (formerly www.deep-er.eu) for this. The main reason for this decision was, that the DEEP-ER website has already a solid audience, plus it has undergone already a major re-launch and fits perfectly the communication needs of the project.

The new website can be reached via www.deep-projects.eu.

In a first step, we have adapted the 2nd level navigation – the first level navigation has proven to be adequate – to integrate the new content on DEEP-EST, e.g. information on the new research fields like resource management and job scheduling, on the new partners, the new applications and the like. In a second step, we have started to incorporate and will continue to do so information on the first project DEEP – information that is at the moment still available via www.deep-project.eu.

2.3.2 Social Media

In the same way, the existing social media channels will be kept alive. For the DEEP and DEEP-ER projects we started a [LinkedIn](#) group and a [Twitter](#) handle. Both channels have been launched under the DEEP projects brand already, hence only slight adaptations were necessary (e.g. updating the header images, slight adaptations to the ‘about’ texts). Whereas we use Twitter more for short news, LinkedIn is a bit more ‘corporate’ e.g. for sharing job posts as well. Also, both channels have a bit different audience. The audience on LinkedIn are e.g. internal project colleagues, and people who know the project and its members personally all to a large degree from an HPC background, the audience on Twitter is more diverse and more global.

To address an even bigger audience, we have decided to start a [Facebook](#) page on top. We hope to be able to reach out to people of different age groups, a more female audience as well as maybe even a more general public.

Future plans:

- **Twitter** has proven to be the most valuable social media channel for the project – user statistics as well as interactions are good and developing as planned, hence, the outreach strategy for this channel will not face major changes.
- **LinkedIn** has been rather slow for some time. For now, we have decided to keep the group alive for some time to come and compare the effectivity of the LinkedIn outreach to the effectivity of the outreach on the new channel Facebook.
- **Facebook** has only been set-up very recently. Since Facebook might give us a better shot at targeting further groups like e.g. more females, maybe a more ‘non-technical’ audience, we plan to adjust the content for this channel as compared to e.g. Twitter. Content will most likely be more high-level, behind the scenes and team stories, more focus on applications. Still, we have to regularly check the metrics on the various posts and will be open to vary and adapt content.

For Twitter and Facebook, we might eventually look into sponsored posts, thus investing into paid content, however, for selected content only. This could be e.g. for promoting our early access programme.

Last but not least, we aim to increase outreach on social media leveraging the channels of both partner institutions and those of individual researchers by motioning them, having them like, retweet, repost etc. our messages.

2.4 News & Media Relations

We intend to secure a steady flow of news about the project. Regular updates for and creation of new content for the project will be key in this respect as this will form the basis for outreach via social media and the e-newsletter. To increase the impact of these communication activities, media relations will be extended and new ones formed. Outreach to the media will be done via classic PR tools like press releases, interviews, background talks etc.

Over the years, the DEEP projects have built a network of media contacts to key outlets especially within the HPC media landscape. These include contacts to e.g.

- [insideHPC](#)
- [Scientific Computing](#)
- [ScienceNode](#)
- [HPCwire](#)
- [The Next Platform.](#)

This list will be extended to include in particular media channels outside the HPC environment.

Relevant media include e.g.:

- [DataCenterInsider](#) (German trade media)
- [Heise online](#) (German trade media)
- Public radio e.g. [Deutschlandfunk.](#)

Additionally, we will look into leveraging news portals frequently used by science journalists like Eurekalert (<https://www.eurekalert.org>). The site is regularly visited by some 8,000 registered journalists from 75 countries and lets you publish on project news, peer-reviewed publications of your project etc.

2.5 Liaison with influencers & European collaboration

Next to traditional media and social media influencers, it will also be key for us to reach out to influencers in HPC predominantly on a European, but also on a global level. A special focus here will be on lobbying groups, special interest groups and standardisation bodies. It is important to mention here, that these kinds of activities are carried out by various DEEP projects partners as many of them are already active members of those groups and will have ample opportunities to promote the project's technologies and developments. WP7 will help to support the partners where necessary and possible e.g. with providing materials. The groups include e.g.:

- **ETP4HPC:** Quite a number of DEEP-EST partners are members or even founders of ETP4HPC and partly also lead working groups within ETP4HPC. Over the timeframe of all DEEP projects the influence has been mutual and on all areas relevant to the project.

- **OpenMP forum:** BSC is an active member of the OpenMP community and pushes OmpSs developments into the OpenMP standard.
- **MPI Forum:** Findings concerning MPI-related extensions to ParaStation will be published by JUELICH's linked third party ParTec to the MPI community, giving feedback to the MPI Forum as the standard setter.
- **SLURM:** BSC and JUELICH will use its already existing contacts within the SLURM community to push the DEEP EST extensions into the SLURM's main branch.
- **Big Data Value Association (BDVA):** Intel, BSC, and FHG are members and will promote the application and system results of DEEP-EST with an eye on influencing standardisation of HPDA use-cases and interfaces.
- **Energy Efficiency HPC Working Group:** BADW-LRZ has been an active member of this working group initiated by LLNL for years and is co-leading one of the sub-groups.

One further important group we want to influence are **early adopters** of our technologies as extending the use of the prototype will allow to give feedback to DEEP-EST partners on the hardware configuration, system software functionality, and general usability of the prototype. More details on what is planned in this respect in section 3.

Finally, collaborations with other European projects and the EC itself have been vital for DEEP projects and will continue to be in the future.

2.6 Dissemination of results: Scientific Publications, Deliverables & Technical Documentation

The project partners will aim at publishing the project's results in peer-reviewed papers and proceedings of relevant scientific conferences. As stressed already in the DoA, the consortium fully supports the open access (OA) strategy put forward by the EC for H2020 research projects and will strive for OA in all publications. Budget has been reserved for this. Speaking of experience, however, a mix between the 'gold' and 'green' open access model seems more feasible to achieve. Coordinating partner JUELICH will make available their repository JuSER (JUELICH Shared Electronic Resources) for publishing scientific articles according to the green model. JuSER is certified by OpenAire.

In a second step, it will be important to make our target groups, in this case especially the scientific ones, aware of the publications and also of the presentations given at conferences and workshops. We will collect and put online links to all peer-reviewed publications on the project website and where possible share for downloading the presentations (<http://www.deep-projects.eu/project/publications.html>). We will also spread the word on the publications via our social media channels and highlight it in the e-newsletter.

2.6.1 List of target publications (not comprehensive)

HPC / IT journals

- Supercomputing Frontiers and Innovations
- Concurrency and Computation: Practice and Experience (Wiley)
- Advances in Engineering Software (Elsevier)
- Transactions on Architecture and Code Optimisation (ACM Digital Library)
- Parallel Computing (Elsevier)
- Journal of Systems Architecture (Elsevier)

- International Journal of Parallel Programming (Springer)
- Journal of Machine Learning Research, Data Mining and Knowledge Discovery (Springer)
- Lecture Notes in Computer Science

Relevant domain science journals

- Astrophysical Journal
- Physics of Plasmas
- Geoscientific model development
- Journal of Computational Physics

Proceedings of relevant conferences

- International Conference on High Performance Computing (HiPC)
- International Conference on High Performance Computing and Simulation (HPCS)
- International Conference on Parallel Processing (ICPP)
- IEEE Cluster
- IEEE International Parallel and Distributed Processing Symposium
- ISC High Performance
- Memsys
- Supercomputing

2.6.2 Deliverables & Technical Documentation

Both the deliverables and the technical documentation contain very valuable project results. As far as the deliverables are concerned, they will be made available via the project website once approved by the EC and depending on their status (public or not). Please note, this link might change. As mentioned, we are still in the process of integrating also information of the first project DEEP into the DEEP projects website. Most probably there will be three lists of deliverables (one for DEEP, one for DEEP-ER and one for DEEP-EST) and the link above might feature links to the three separate lists.

With respect to the technical documentation: Especially the academic partners work with and develop further open source software by their institutions and publish these on their institutions' website, GitHub or similar. Hence, disseminating it, will most likely be a collection of links in a dedicated space on the website (currently here: <http://www.deep-projects.eu/software/sources.html>).

2.7 Conferences, events & trade fairs

Next to medially transported communication (via owned or earned media) it is crucial for DEEP-EST's communication strategy to foster one-to-one encounters with our key target groups at conferences, trade fairs and public events.

With respect to scientific conferences, possible activities include participation in the technical programmes via BoF-sessions, workshops, tutorials or research papers and posters. We will also explore the possibility to exhibit on the show floors. This will be dependent, however, on e.g. potential collaborations with other projects or on the activities in this area by consortium partners to keep costs and benefits in an acceptable ratio.

We will also engage in self-organised events like the planned user awareness days or e.g. roundtables as well as present the project to the public at e.g. open door days at the partners' premises.

2.7.1 *List of target conferences & events (not comprehensive)*

HPC / IT events and conferences

- European HPC Summit Week
- ISC High Performance
- Supercomputing

Industry-related and domain science events and conferences

- CeBIT
- Hannover Messe
- AGU & EGU (American and European Geoscience Union)
- Environmental Informatics
- IEEE ICMLA (IEEE International Conference on Machine Learning and Applications)

3 Early Access Programme

One of the challenges for the DEEP-EST project is to ensure we produce a prototype system that can support, and significantly improve performance for the widest range of applications and workloads possible in computational simulation and data analytics.

In the project, we are focussing on a broad range of applications, represented by our six co-design applications. However, we need to bring in a wider range of applications and users to allow the computational simulation and data analytic communities to exploit the benefits of our hardware and software platform and to allow us to ensure that our systems support the wide range of application and usage patterns required for a general-purpose system.

Extending the use of the prototype will allow feedback to DEEP-EST partners on the hardware configuration, system software functionality, and general usability of the prototype. This will facilitate improved design and inform any future production systems. Furthermore, it will enable user communities to evaluate any changes required for their applications or workflows, widening the co-design impact of the project and ensure applications are able to use the new generation of HPC and HPDA systems that will be based on the DEEP-EST prototype and software.

We are currently designing our early access programme to ensure we reach a wide range of potential users, both from academia and industry. We will target specific industry and academic events; i.e. conferences and workshops, with promotional literature and talks to inform potential users about the early access programme.

We will setup an early access website and mailing list, to provide data and communication on the early access programme. Once we have started the dissemination strategy, and we have the prototype ready for early access users, we will launch and run a number of calls for users, where interested parties submit a brief application, including some information on the applications that would be used.

We will establish an early access programme panel, which will be composed of project partners and/or external experts (as required), who will assess the applications and choose the applications that will get the best benefit from early access to the prototype, and will provide the broadest range of coverage and feedback to DEEP-EST.

As one part of the deal, early access users will have to commit to reference the DEEP-EST project when they talk about their experiences or even publish on them. Also, we would like them to share their use cases with us that we can collect in a case study brochure together with use cases from the co-design applications from WP1.

4 Training & Education

Finally, education and training are key for the DEEP projects outreach and exploitation strategy. In close cooperation with WP1 Applications, a systematic education and training programme will be implemented.

The goal of the programme is twofold:

- **Internal:** Between the partners, targeted training sessions and workshops will ensure the exchange of knowledge and background relevant to the R&D activities of the consortium. Most training efforts will be directed at application developers. Despite the target group being internal, this will have relevance to external target groups as well, as we intend to make most of the material available via the project website.
- **External:** Outside of the project, the consortium will train stakeholders specific to their interests in DEEP-EST developments and their needs. In a first step, it will be crucial to train interested users as part of the early access programme. But efforts will not be limited to this target group. Taking part in workshops at events and conferences will also be sought after.

In the past, co-organisation of training courses and workshops with other related projects has proven to be of mutual benefit to the participants. The DEEP projects will seek collaborations like this going forward.

4.1 Training topics

It is important to mention that internal trainings and workshops for all partners will depend heavily on the need for further education and might vary – also depending on the previous knowledge especially of the application partners new to the consortium. This is for instance the case with Intel Xeon Phi trainings. Since quite elaborate training material already exists on KNL (e.g. the) and application developers – both DEEP project veterans and new partners – have already gained far-reaching experience, intensive Xeon Phi trainings might not be necessary. For topics like e.g. NVM, NAM or GCE, the co-design discussions will lead to concrete use cases for these innovative technologies. Training session will be organised if needed and closely aligned to the use cases to be tested within the timeframe of the project.

It is foreseeable, however, that internal trainings and workshops in DEEP-EST will most likely cover the following topics.

- **Performance and benchmarking tools:** One important point at the beginning of the project is to analyse the applications in detail. This is crucial to understand how they can benefit the most from the modular architecture in DEEP-EST. BSC provides

several tools for analysing and profiling. For most of the application partners these tools are new, so they will need an introduction and training on how to use the tools to get the information they need out of their applications and deliver valuable input for the co-design process.

- **OmpSs:** The programming model developed by BSC is new to some of the application partners in DEEP-EST. Introducing those partners to OmpSs and its specific extensions for a modular supercomputing architecture – especially the Cluster-Booster partition and the OmpSs offload functions – will be necessary. On top, it will be beneficial to keep the partners that are already familiar with OmpSs updated on the latest OmpSs features.
- **DEEP projects I/O stack & resiliency features:** In the DEEP-ER project novel I/O concepts and improved resiliency features (application and task-based checkpoint-restart functions) have been developed to exploit the advancements of a multi-layered memory hierarchy. Especially application partners new to the project will benefit from trainings on this.
- **Data Analytics:** Among the application partners in WP1 a group has already formed that will look more detailed into data analytics. Again, depending on their need for training, Data Analytics courses might be scheduled within the DEEP-EST project or partners will be able to attend courses offered by PRACE. Also, partners JUELICH and BADW-LRZ are both expanding their courses in the fields of Deep Learning and Machine Learning.

4.2 Training Schedule

In the beginning, a major focus of the internal training will be on familiarising the application partners with the performance analysis and benchmarking tools brought to the project by BSC (e.g. Extrae and Paraver) and by JUELICH (JUBE). Planning has already started to organise a 2-days workshop at BSC to be held from 28 to 29 November 2017. On top, Xeon Phi trainings might be organised, if needed. Also, first introductions to OmpSs (with a focus on optimisations for the Booster) form part of the training efforts in the first half of the project.

In the second half, trainings on I/O and resiliency features are more in the focus. Special training sessions closely aligned to possible use cases will take centre stage during the final phase of the project.

In addition, all DEEP-EST consortium members will have the possibility to attend PRACE PATC courses on topics of interest as BADW-LRZ, BSC, EPCC, JUELICH and NCSA are all PRACE training centres. Since PRACE is intensifying its offer of courses on data science in the future, the DEEP-EST partners will have a choice of several courses – and even a MOOC (Massive Open Online Course) – on data analytics which will enable them to cope with the new DEEP-EST Data Analytics Module.

As far as external workshops are concerned, DEEP-EST colleagues have already been and are promoting the general idea of the Modular Supercomputing Architecture. Towards the middle and the end of the project, external workshop and education efforts will focus more on details of the R&D developments.

4.3 Course material

Most importantly, all course material will be made available to project partners through the BSCW document sharing platform (in the directory: /DEEP-EST/Dissemination and Training/Training). Where legally possible, the material will also be made publicly available via the project website, as we have started to do so in DEEP-ER (<http://www.deep-projects.eu/applications/trainings.html>).

5 Outlook

With the communication and outreach activities in DEEP and DEEP-ER we have laid a good basis: We can draw on e.g. a database of already 300 contacts, a professional looking website, Twitter as most beneficial social media channel for the DEEP projects, a network of collaborating EU projects, good contacts to e.g. standardisation bodies and lobbying groups. The first and foremost goal is to keep up this work and improve and extend it in the areas described in the document. In the first six to 12 months a focus will be on:

- Spreading the word on the MSA.
- Fully integrating all DEEP projects content into the website.
- Bring up to speed the Facebook channel and the e-newsletter campaign.
- Flesh out the conditions of the early access programme.
- Start the training programme.

Annex A

A.1 Communication Toolkit for Download

Flyer DEEP-EST

Flyer DEEP Projects

Project logos

List of Acronyms and Abbreviations

A

ASTRON: Netherlands Institute for Radio Astronomy, Netherlands

B

BADW-LRZ: Leibniz-Rechenzentrum der Bayerischen Akademie der Wissenschaften. Computing Centre, Garching, Germany

BDA: Big Data Analytics

BDEC: Big Data and Extreme-Scale Computing

BeeGFS: The Fraunhofer Parallel Cluster File System (previously acronym FhGFS). A high-performance parallel file system.

BN: Booster Node (functional entity)

BoP: Board of Partners for the DEEP-EST project

BSC: Barcelona Supercomputing Centre, Spain

BSCW: Repository used in the DEEP-EST project to share all project documentation.

C

CA: Consortium Agreement

CD: Corporate Design

CERN: European Organisation for Nuclear Research / Organisation Européenne pour la Recherche Nucléaire, International organisation

CM: Cluster Module: with its Cluster Nodes (CN) containing high-end general-purpose processors and a relatively large amount of memory per core

CN: Cluster Node (functional entity)

CPU: Central Processing Unit

D

DAM: Data Analytics Module: with nodes (DN) based on general-purpose processors, a huge amount of (non-volatile) memory per core, and support for the specific requirements of data-intensive applications

DDG: Design and Developer Group of the DEEP-EST project

DEEP: Dynamical Exascale Entry Platform (project FP7-ICT-287530)

DEEP-ER: DEEP - Extended Reach (project FP7-ICT-610476)

DEEP/-ER:	Term used to refer jointly to the DEEP and DEEP-ER projects
DEEP-EST:	DEEP - Extreme Scale Technologies
Dimemas:	Performance analysis tool developed by BSC
DN:	Nodes of the DAM
DoA:	Description of Action

E

EC:	European Commission
EEHPC:	Energy Efficient High Performance Computing
EEP:	European Exascale Projects
EPT4HPC:	European Technology Platform for High Performance Computing
ESB:	Extreme Scale Booster: with highly energy-efficient many-core processors as Booster Nodes (BN), but a reduced amount of memory per core at high bandwidth
EU:	European Union
Exascale:	Computer systems or Applications, which are able to run with a performance above 10^{18} Floating point operations per second
EXDCI:	European Extreme Data & Computing Initiative
EXTOLL:	High speed interconnect technology for HPC developed by UHEI
Extrae:	Performance analysis tool developed by BSC

F

FHG-ITWM:	Fraunhofer Gesellschaft zur Foerderung der Angewandten Forschungs e.V., Germany
Flop/s:	Floating point Operation per second
FP7:	European Commission 7 th Framework Programme
FPGA:	Field-Programmable Gate Array, Integrated circuit to be configured by the customer or designer after manufacturing

G

GFlop/s:	Gigaflop, 10^9 Floating point operations per second
GROMACS:	A toolbox for molecular dynamics calculations providing a rich set of calculation types, preparation and analysis tools

H

H2020:	Horizon 2020
HBM:	High Bandwidth Memory
HPC:	High Performance Computing
HPDA:	High Performance Data Analytics
HW:	Hardware

I

IC:	Innovative Council
Intel:	Intel Germany GmbH, Neubiberg, Germany
I/O:	Input/Output. May describe the respective logical function of a computer system or a certain physical instantiation

J

JUBE:	Jülich Benchmarking Environment
JUELICH:	Forschungszentrum Jülich GmbH, Jülich, Germany
JURECA:	Jülich Research on Exascale Cluster Architectures
JuSER:	JUELICH Shared Electronic Resources (online repository)

K

KNL:	Knights Landing, second generation of Intel® Xeon Phi™
KPI:	Key Performance Indicator
KU Leuven:	Katholieke Universiteit Leuven, Belgium

L

LHC:	Large Hadron Collider (LHC), the world's most powerful accelerator providing research facilities for High Energy Physics researchers across the globe
-------------	---

M

Megware:	Megware Computer Vertrieb und Service GmbH, Chemnitz, Germany
MHD:	Magneto-hydrodynamics

Mont-Blanc:	European scalable and power efficient HPC platform based on low-power embedded technology
MoU:	Memorandum of Understanding
MPI:	Message Passing Interface, API specification typically used in parallel programmes that allows processes to communicate with one another by sending and receiving messages
MSA:	Modular Supercomputer Architecture

N

NAM:	Network Attached Memory
NCSA:	National Centre for Supercomputing Applications, Bulgaria
NEST:	Widely-used, publically available simulation software for spiking neural network models developed by NMBU.
NF:	Network Federation within the DEEP-EST prototype
NMBU:	Norwegian University of Life Sciences, Norway
NVM:	Non-Volatile Memory. Used to describe a physical technology or the use of such technology in a non-block-oriented way in a computer system
NVRAM:	Non-Volatile Random-Access Memory

O

OA:	Open Access
ODC:	Other direct costs
OmpSs:	BSC's Superscalar (Ss) for OpenMP
OpenCL:	Open Computing Language, framework for writing programmes that execute across heterogeneous platforms
openHPC:	A community effort that is initiated from a desire to aggregate a number of common ingredients required to deploy and manage HPC Linux clusters
OpenMP:	Open Multi-Processing, Application programming interface that support multiplatform shared memory multiprocessing

P

ParaStation:	Software for cluster management and control developed by JUELICH and its linked third party ParTec
Paraver:	Performance analysis tool developed by BSC

ParTec:	ParTec Cluster Competence Center GmbH, Munich, Germany. Linked third Party of JUELICH in DEEP-EST
PFlop/s:	Petaflop, 10^{15} Floating point operations per second
PI:	Principal Investigator
PMT:	Project Management Team of the DEEP-EST project
PRACE:	Partnership for Advanced Computing in Europe (EU project, European HPC infrastructure)

Q

R

R&D:	Research and Development
RDA:	Research Data Alliance
RM:	Resource Manager

S

SCR:	Scalable Checkpoint/Restart. A library from LLNL
SDV:	Software Development Vehicle: HW systems to develop software in the time frame where the DEEP-EEST prototype is not yet available.
SIONlib:	Parallel I/O library developed by Forschungszentrum Jülich
SLURM:	Job scheduler that will be used and extended in the DEEP-EST prototype
Slurm:	MHD code developed by KULeuven.
SME:	Small and Medium Enterprises
SRA:	Strategic Research Agenda prepared by ETP4HPC
SSSM:	Scalable Storage Service Module
STEM:	Science, technology, engineering and mathematics
SW:	Software

T

TFlops:	Teraflop, 10^{12} Floating point operations per second
Tk:	Task, Followed by a number, term to designate a Task inside a Work Package of the DEEP-EST project
ToW:	Team of Work Package leaders of the DEEP-EST project

U

UEDIN:	University of Edinburgh, UK
UHEI:	Ruprecht-Karls-Universitaet Heidelberg, Germany
Uoi:	Háskóli Íslands – University of Iceland, Iceland
UPC:	Universitat Politècnica de Catalunya. Barcelona, Spain

V

W

WP:	Work package
------------	--------------

X

Y

Z