Deliverable 1.2
Results of a Survey concerning the Tier-2 and Tier-3 HPC-Infrastructure in Germany

Work Package 1

August 2017

ProfiT-HPC Consortium

Gefördert durch die Deutsche Forschungsgemeinschaft (DFG)
Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)
KO 3394/14-1, OL 241/3-1, RE 1389/9-1, VO 1262/1-1, YA 191/10-1

Abstract

In this project, the goal is to aid especially new users with only little programming experience in understanding the importance of performance profiling and optimisation of an application. For this, we will develop a toolkit which shall provide HPC users with an understandable report, at best in a comprehensive, consistent manner throughout Germany by rolling out the toolkit during the course of the project.

To gain some insight of the hardware and software infrastructure at Tier-2 and 3 systems in Germany, a survey was created and disseminated to 53 institutions. The dissemination was limited to computing centres and institutes at universities with compute clusters - no companies were invited to answer the survey, as these are not the target audience of the project. The result of the survey will give an idea, which concepts can be followed and which concepts can be neglected, because they are not available or not supported on site.

This deliverable presents the results of the survey and especially explains the (possible) impact of the survey results on the project. We would like to thank all of the institutions that have taken part in the survey and hope that the gathered information can also help other projects and developers.

Contents

Gefördert durch die Deutsche Forschungsgemeinschaft (DFG)
KO 3394/14-1, OL 241/3-1, RE 1389/9-1, VO 1262/1-1, YA 191/10-1

Introduction and Methodology

This document will describe the Tier-2 and Tier-3 high performance computing centres infrastructure in Germany, based on the survey designed and created as part of task 1.2 in the DFG project "ProfiT-HPC - Profiling Toolkit for HPC in Tier-2 and 3". It will give an overview of the basic infrastructure components of Tier-2 and Tier-3 computing centres in terms of hardware (CPUs, interconnects, parallel file systems, accelerators, etc.) and software (operating systems, batch systems, C/C++ and Fortran compiler families, MPI libraries, profiling tools, etc.) and shall serve as a profound basis for the development of the toolkit. This information is needed, for example, to evaluate whether the selected tools which are described in D2.1 [**pfit2017-2.1**] and D2.2.1 [**pfit2017-2.2.1**], are installed on site. Each additionally necessary tool that needs to be installed on the systems might pose a great hurdle in the roll-out of the toolkit at the end of the project. In addition, it is important to know which metric collection tools are already available at the institutions. If there is a globally available metric collection tool, it will make sense to use this as a basis for the developed toolkit.

The information was gathered by an online survey created with and provided through the LimeSurvey tool at the Universität Hamburg[1]. The survey consisted of 7 categories (hardware, operating systems, batch systems, C/C++ and Fortran compiler families, parallel programming tools and libraries, profiling and monitoring tools, scientific application software and languages) and 17 questions, all of which are listed in the appendix (see p. 15).

The invitations were sent out step-wise by e-mail, beginning with the HLRN computing centres, which was followed by 3 other invitations steps. The survey started in the beginning of April 2017 and ended on July 31th, 2017. In total, 53 institutions were invited to complete the survey, of which 33 followed the invitation and provided information for a total of 34 systems, which is a response from about 62 % of all invited computing centres. In the following, the results of this survey will be described and visualized where appropriate. In all cases the results will be put into the context of our project by explaining the possible impact.

[1] https://www.rrz.uni-hamburg.de/services/software/alphabetisch/limesurvey.html

Infrastructure Survey Results

The following 33 computing centres and universities/institutes have participated at the survey (53 institutions were contacted):

- ZIB (Zuse Institute Berlin; for HLRN),
- GWDG (Gesellschaft für wissenschaftliche Datenverarbeitung mbH Göttingen),
- Universität Oldenburg (ForWind),
- Technische Universität Hamburg-Harburg,
- Zentrum für angewandte Raumfahrttechnologie und Mikrogravitation (ZARM), Universität Bremen,
- Jacobs University Bremen, Computational Laboratory for Analysis, Modeling and Visualization,
- Rechenzentrum, Christian-Albrechts Universität zu Kiel,
- Regionales Rechenzentrum der Universität Hamburg,
- Regionales Rechenzentrum Erlangen (RRZE), Universität Erlangen,
- Kommunikations- und Informationszentrum (kiz), Universität Ulm,
- Paderborn Center for Parallel Computing, Universität Paderborn,
- Regionales Hochschulrechenzentrum Kaiserslautern (RHRK), TU Kaiserslautern,
- IT Center, RWTH Aachen,
- Helmholtz-Zentrum Dresden-Rossendorf,
- Potsdam-Institut für Klimafolgenforschung,
- Hochschulrechenzentrum (HRZ), Universität Marburg,
- Universität Bayreuth, IT-Servicezentrum,
- BTU Cottbus - Senftenberg (CFTM2),
- Universität Kassel,
- Universität Bielefeld, Fakultät für Physik (Theoretische Hochenergiephysik),
- Alfred-Wegener-Institut, Helmholtz-Zentrum für Polar- und Meeresforschung,
- Universität Rostock (ITMZ),
- Deutsches Klimarechenzentrum GmbH,
- Leibniz Universität IT Services (LUIS, formerly RRZN),
- Technische Universität Dortmund,
- IT-Servicezentrum, Martin-Luther-Universität Halle-Wittenberg,
- TU Bergakademie Freiberg,
- Otto-von-Guericke-Universität Magdeburg,
- Technische Universität Darmstadt (Hochschulrechenzentrum),
- Zentrum für Datenverarbeitung, Eberhard Karls Universität Tübingen,
- URZ Heidelberg/RUM Mannheim,
- and two other institutions, which could not be identified.

Each of these institutions have answered the questions listed in App. A, the results of which will be presented in the following subsections. The structure will closely follow the structure of the survey, making small changes in the ordering only where this aids in putting the results into a context. Some results shown in this section omit certain answers or combine single answers into groupings for the sake of readability and according to the relevance to our project. The raw data can also be found in App. A.

Hardware

Processor Types

In this subsection the installed CPU types are listed and the count of each CPU type is given. The computing centres named all CPUs installed in the system, thus the count of processor types is much higher than the number of participants. Different processors have, for example, different sets of hardware event counters which the performance monitoring or profiling tools can access. Not only the actual counters are different, but also the naming conventions differ. Thus, an overview of the different types of processors will give us a minimal set of hardware performance counters to rely on.

The Intel table on p. 15 in the appendix lists all Intel processors that the compute centres participating in the survey have installed on their sites. In total, 86 CPUs (with 48 different Intel Xeon processor types) are present in the clusters, of which the top three are:

1. Intel Xeon E5-2680 v3 (8),
2. Intel Xeon E5-2670 (7), and
3. Intel Xeon E5-2630 v3 (5).

As can be seen in the table on p. 15, different version numbers of the same processor are built into the clusters. The version numbers resemble the generation of each of the processors and accordingly the different architectures. Since the architectures are the deciding factor in terms of hardware counters, the following plot shows the CPU distribution with respect to their generation (Broadwell, Haswell, etc.):

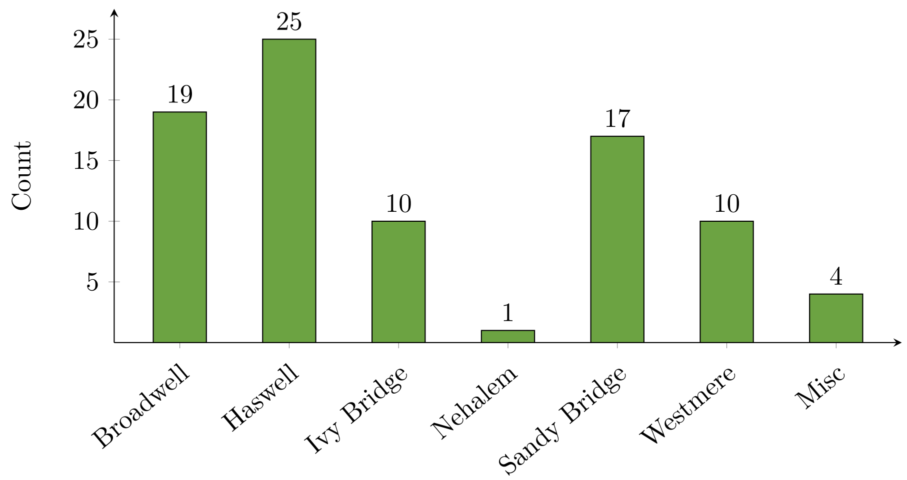

About 29 percent of the Intel CPUs belong to the Haswell generation, followed by Broadwell (22 %), Sandy Bridge (20 %) and Ivy Bridge and Westmere (each 12 %). There are no newer CPU generations of the Core i series (for example Skylake), but this may change during the course of the project.

In addition to the clusters with Intel Xeon CPUs, there are also clusters with AMD CPUs. These are listed in Tab. 2 in the appendix. Nine different AMD processor types were named, all from the Opteron series. Overall, the number of Intel processors integrated in the systems is much higher than that of AMD processors (86 Intel vs. 12 AMD CPUs). Thus, in this project we will first concentrate on Intel architectures as far as applicable.

Accelerators

This subsection considers the accelerator capacities of the institutions. As before, multiple selections were possible.

Here, the Nvidia accelerators are presented, while the Intel Xeon Phi accelerators can be found in the appendix because of their minor importance compared to the Nvidia accelerators. No institution uses AMD accelerator cards on the nodes. The complete answers are shown in tables on p. 16 and 17.

Due to the wide spectrum of available Nvidia accelerators, only the architecture types of the graphic cards are shown in the bar chart to get an overview of each architecture generation:

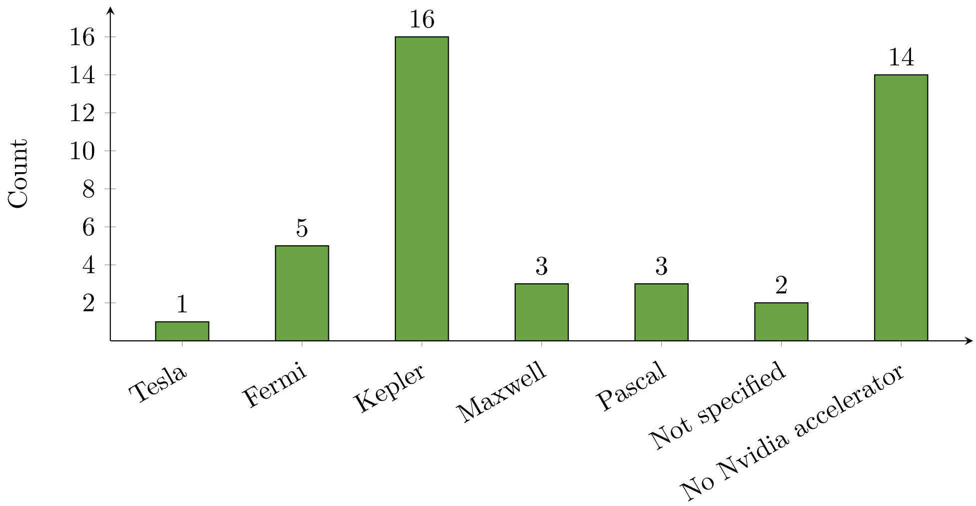

Considering the results in the appendix, we can see that the accelerator cards by Nvidia (especially the Kepler architecture) are dominating the area of accelerators in a ratio of 30:7 compared to the Intel Xeon Phi co-processors (which are installed at six compute centres). Therefore, the Intel co-processors are not included here.

Interconnects

This subsection deals with the different types of the communication network families used in the clusters. It is possible that more than one communication network type was marked by one institution, because multiple choices were allowed. The information on used interconnects is useful because it might be possible to use tools which are tailored to a special kind of interconnect.

The explicit answers are listed on p. 16 and can be grouped into the following families:

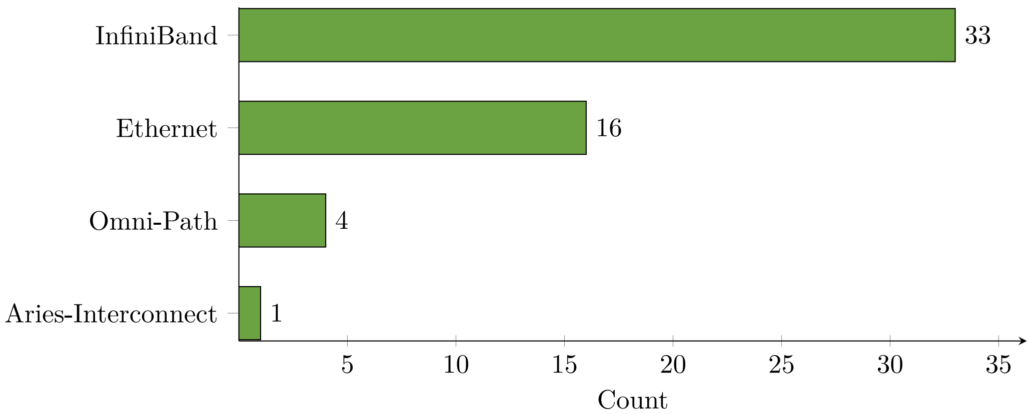

The main three interconnects are QDR-Infiniband (19), Gigabit Ethernet (14) and FDR-Infiniband (10). The main usage of Ethernet interconnects will most probably be the connection to the file system, but nonetheless, it is to be considered. Surprisingly, already four centres use Omni-Path as their interconnect.

File Systems and Memory

In the following table, the count of the different parallel file systems is given. Again, multiple selections were possible.

Parallel filesystem	Number of answers
BeeGFS	12
GPFS	9
Lustre	8
Panasas	2
ScaTeFS	1
No answer	3

It can be seen, that the freely available BeeGFS file system is the most prevalent parallel file system, followed by the proprietary GPFS and Lustre. This will be interesting for the collection of file I/O data.

In addition to the used parallel file systems, also the existence of non-POSIX memory systems (like object stores) on the clusters is of interest. This was a yes or no answer showing that there are almost no non-POSIX memory systems on site (with 2 exceptions), as can be seen on p. 17 in the appendix.

Operating Systems

In the survey section *Operating Systems*, questions about which operating systems are running on the clusters and which distribution and kernel version is used were asked. Furthermore, information was gathered on proprietary distributions and whether cgroups and virtualisation techniques are used. This very specific and detailed information can again be helpful for information on the kernel-specific availability of certain tools like perf [**gregg_perf**], the possibility of using cgroups for accounting resource usage and whether virtualisation techniques play an important role and thus need to be considered in the toolkit. The results of these questions are discussed in the following subsections.

Different Operating Systems, Distributions and Kernels

The first question concerned the operating systems running on each cluster, the distribution and the kernel version used. This information is needed to know on which operating system platform the profiling toolkit should be implemented and which kernel features can be used. This was a free text field answer and more than one answer were possible. To no surprise, Linux is the only operating system on all clusters.

The complete list of distributions in combination with kernel version numbers can be found in the appendix, while here the most important results are summarised. The following chart thus only shows the different distributions without version numbers. For better readability, SUSE Linux Enterprise Server (SLES) and Red Hat Enterprise Linuux Server (RHEL) are abbreviated in the chart.

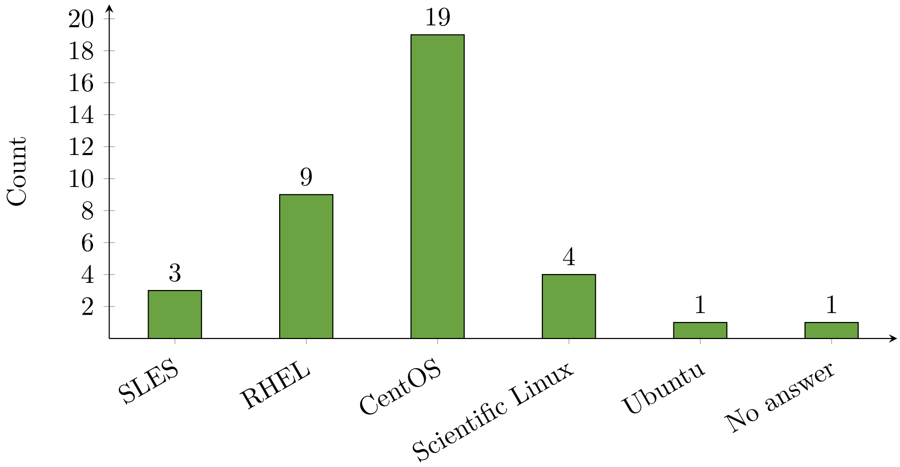

We can conclude, that the Red Hat Enterprise Linux Server distribution and the derivatives of RHEL are dominating the national distribution market for HPC very strongly, whereas SUSE and Ubuntu play only a minor role.

The different kernel versions on site were also of interest, because some tools need special kernel versions. All kernel versions used are listed on p. 17. Only 23 of the 33 institutions gave information on the kernel version used. The resulting answers show that most clusters are using kernel version 3.10 or other minor versions of the major kernel version number 3. Even though there is a clear peak at kernel version 3.10. A total of nine different kernel versions are in use, indicating a very heterogeneous kernel updating behaviour across the different compute centres.

Proprietary Linux Environments

A proprietary Linux environment comes with additional benefits like support, which could potentially be exploited by the toolkit. The answer input was again entered via a text field and is listed below:

Proprietary Linux	Number of answers
Bull Supercomputing Suite	1
Cray Linux Environment	1
ParTec ParaStation	2
Red Hat Enterprise Linux Server	1
No	28

It can be seen from the table, that most of the institutions are using free Linux environments.

Virtualisation Techniques

This subsection treats the results of the question, whether container or virtualisation techniques are implemented. This information could be necessary to evaluate the values of user or system CPU time. In addition, this will influence the monitoring and profiling options given in general.

Virtualisation used	Number of answers
Yes	1
No	26
No answer	7

It can been seen, that at least one, but at most eight institutions use this technology. Thus, about three quarters of the institutions do not use any container or virtualisation technique. This is not very surprising at this point in time, due to the security issues that arise through, e.g., the usage of Docker. Nonetheless this is a field which is rapidly developing and will need to be kept in mind.

Control Groups

The Linux kernel feature *control groups* (cgroups) allows the administrator to limit, account for, and isolate the resource usage (CPU, memory, disk I/O, network, etc.) of a collection of processes (see quasi citation [**wiki__cgroups**]; see further details in [**cgroups**]). Especially it is possible to account those groups of processes for selected metrics, so that we could potentially get runtime data from the operating system with this feature. The following table shows the result of this question:

cgroups	Number of answers
Yes	11
No	17
No answer	6

It can be seen that only approximately one third - at most one half - of the participating institutes are using cgroups, so we can not rely on this feature of the kernel for our toolkit. There are also problems with handling cgroups on some sites, so this point has to be discussed in detail.

Batch Systems and Usage of the Nodes

Batch Systems

Every cluster needs a batch system, which organises the reservation and allocation of the desired number of nodes, the monitoring and cleaning up after job termination. For the purposes of this project it is interesting, which batch system is installed on site, since extracting system data by the batch system is a possibility to get interesting metrics and basic job data. Since every batch system has its own mechanisms, it is important to know the distribution of those systems. This question was a multiple choice question.

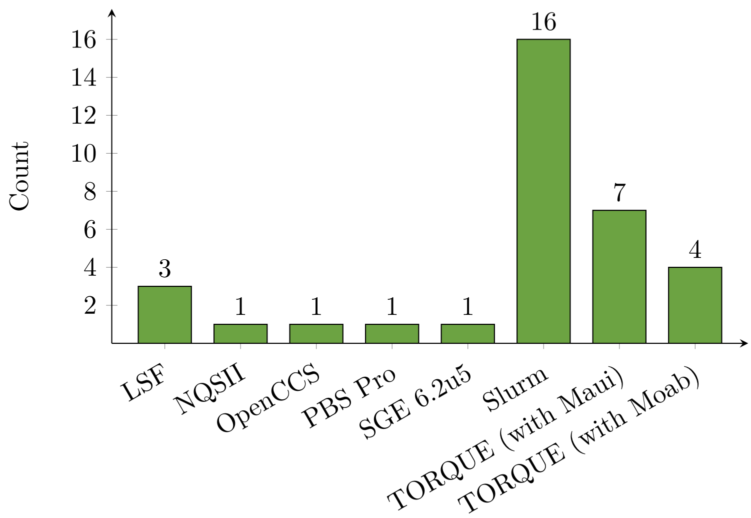

It can bee seen, that the most popular batch systems are Slurm and TORQUE. Both are freely available (the job scheduler Maui too, while the job scheduler Moab is proprietary). Because all of these batch systems are available to the project partners, it is possible to test the mechanisms to get the desired data and learn about other important characteristics.

Usage of the Nodes

The answers to this question can help to identify the usable tools and how they need to be adapted. Many tools are not capable of collecting job-specific metrics but rather collect node-specific data. If nodes are shared by users, a mechanism for collecting job-specific data will have to be implemented. As many institutions have more than one cluster in use, multiple answers were possible.

Node Usage	Number of answers
Non-Shared	14
Shared	27
No answer	0

All institutions have answered this question, and only 14 of 34 clusters have non-shared nodes, i.e., approximately 41 % of the participants have clusters or at least queues with non-shared usage. Quite in contrast, more than 79 % of the institutions also have clusters or queues with shared node usage in operation, an important fact for job-based monitoring, as previously described.

(Parallel) Program Development

C and Fortran Compilers

The question regarding installed compilers arises from the same motivation as the question regarding the batch systems: different compilers offer different flags and options to directly debug or even profile the compiled application. It is a conceivable scenario to add modules or aliases which directly set these flags to gain easy access to application related profiles.

While the complete set of answers can again be found in the appendix, on p. 18, the following bar chart shows the most prevalent compiler families:

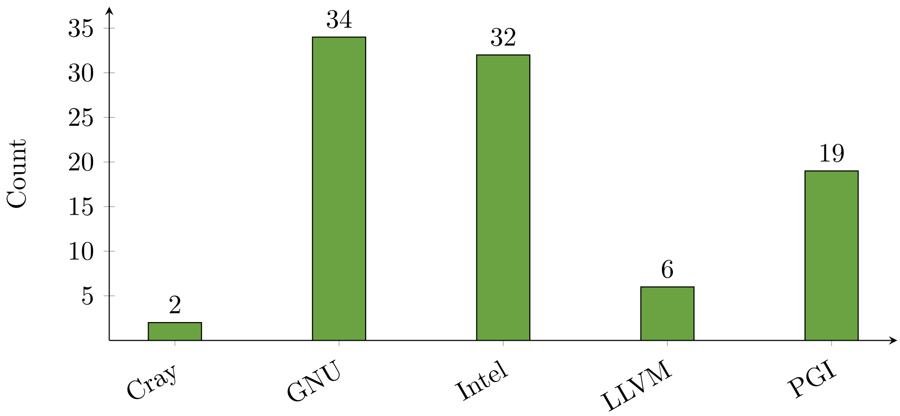

All participants of the survey have the GNU compiler family installed on site and almost all additionally offer the Intel compiler family (32/34). The third rank is occupied by the PGI compiler family, which is located at 19 institutions.

Intel Parallel Studio XE

The specification, whether Intel Parallel Studio XE is installed and if so, which edition is installed, could help to identify Intel tools, which can be used to collect data. For example, the Cluster Edition contains Intel MPI, which can be used to collect network data.

Edition of Parallel Studio	Number of answers
Composer Edition	7
Professional Edition	0
Cluster Edition	17
No	4
No answer	6

Considering these results, there is no basis to include Intel Parallel Studio XE support in our toolkit at the moment.

Installed MPI Implementations

As previously mentioned, different versions of certain soft- and hardware components have different features, some of which might be usable within the toolkit. Intel MPI for example has the ability to collect network traffic data and displays it in a report after the job finished. This feature can be controlled with environment variables and could thus be easily included in a (partially) automated work-flow. As a consequence, the knowledge of which MPI implementations are on site is important to potentially use their special features.

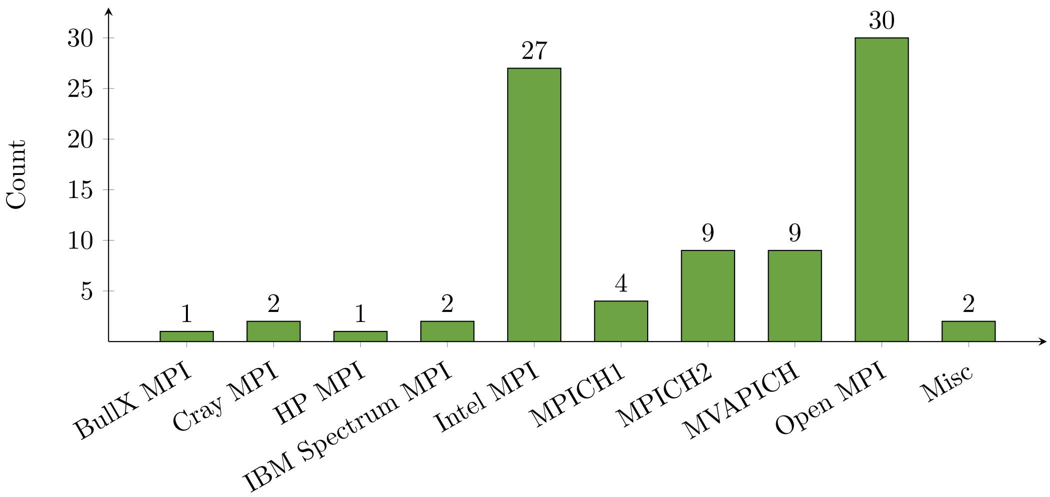

It can be seen, that in most cases Open MPI and Intel MPI are available. With a big gap MPICH2 and MVAPICH are following. As a consequence, it is not possible to rely on any MPI-specific features for an automated toolkit. Even though Intel MPI is offered by many compute centres, the additional availability of other MPI implementations, especially Open MPI, makes it impossible to concentrate on only one implementation. But it is for example feasible to include suggestions for the user to use the above mentioned Intel features or to include mechanisms to take advantage of the features of multiple implementations.

Profiling and Monitoring Tools

The answers in this subsection are crucial for the project, since the goal is to create a lightweight tool-set, based on or including other tools to get performance critical information of an application. For better readability, the Slurm Profile Accounting Plugin API is denoted as Slurm PAP.

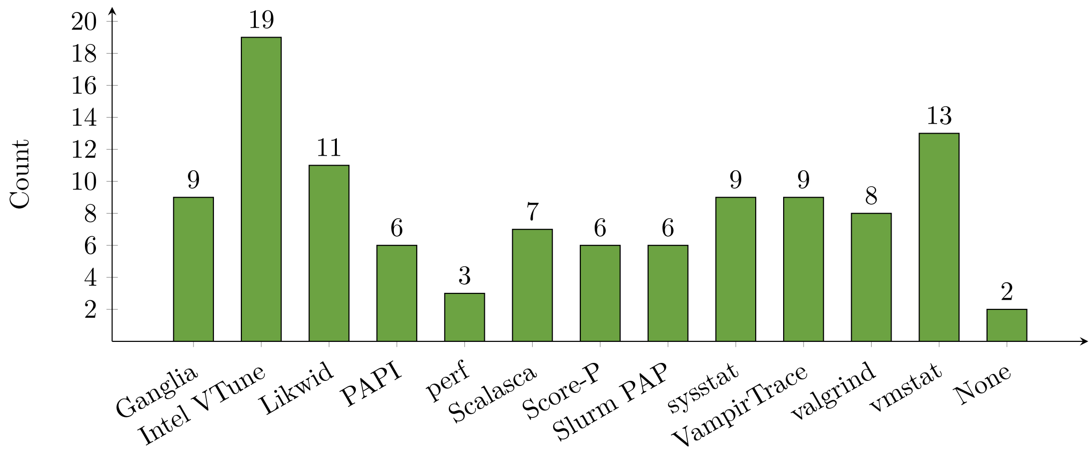

The wide range of available tools is clearly reflected in the answers to this question. Many different tools are installed, none of which on the vast majority of the systems. Intel VTune is clearly the most often installed tool for profiling, followed by vmstat, Likwid, Ganglia, sysstat and VampirTrace. Comparing the list of tools gained from this survey to the list of tools previously evaluated in the scope of this project [**pfit2017-2.1**], we see *Likwid*, *perf*, *sysstat* and *vmstat* to appear on both lists. *Darshan* on the other hand, which would be a good candidate for I/O related performance analysis is not installed on any system (see table on p. 19). This can lead to the choice of a different tool for our final toolkit.

Scientific Application Software and Languages

In this last subsection, the most repeatedly mentioned scientific software packages are listed, to evaluate for which of those packages and programs best practices can be developed. The table on p. 20 lists further programs that appeared in the answers of the survey.

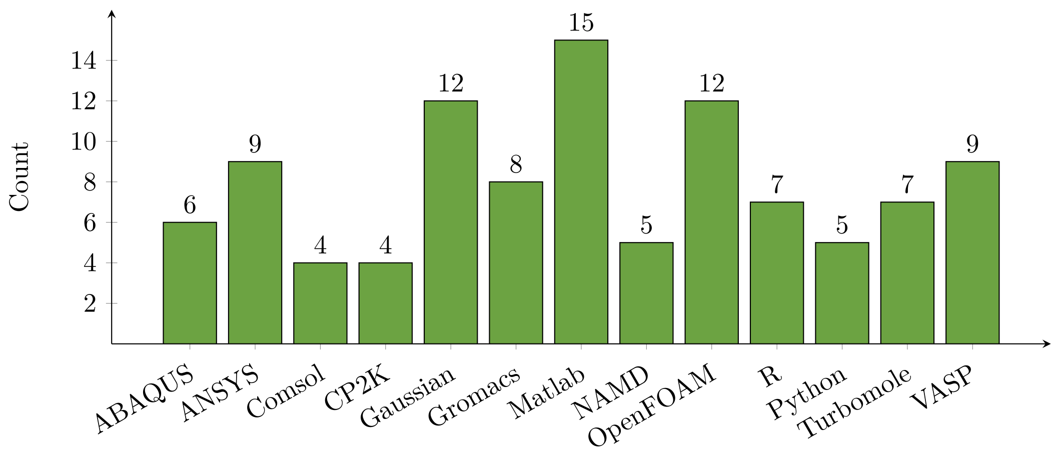

As expected, there is a wide variety of applications used and installed on the systems. The most often mentioned programs are Matlab (15), OpenFOAM (numerical continuum mechanics, 12), Gaussian (computer chemistry, 12), followed by VASP (quantum chemistry, 9), ANSYS (9) and Gromacs (molecular dynamics, 8).

Summary of Survey Results

The first, rather pleasant result of this survey was the high number of participating institutions. Approximately 62 % of the invited institutions took part and answered the questions in detail. In the context of some discussions with participants of the survey, we can state that many cluster providers are highly interested in profiling and there are many cooperation possibilities.

In the following the main results of the survey are again presented. In the field of CPUs it is obvious that Intel dominates the high performance computing world with its Intel Xeon CPUs, while AMD plays only a minor role. Furthermore, from the above results, it is visible that from the 86 installed Intel CPUs there are 54 which were released in the last 5 years (19 Broadwell, 25 Haswell and 10 Ivy Brigde CPUs), while the other 32 are older (including 17 Sandy Bridge CPUs and 11 Westmere/Nehalem architectures).

In terms of the interconnects, a clear picture presents itself since the InfiniBand interconnects dominate (30 from 33 institutions). Here, the QDR InfiniBand is leading (19) followed by the FDR (10). Contrasting, in the field of the parallel file systems there is not such a distinct result because BeeGFS, GPFS and Lustre have approximately (in the context of this project) similar counts (BeeGFS with 12, GPFS 9 and Lustre 8).

Accelerator cards are on site in about 58 % of the clusters and the most used type of graphic card architectures are Nvidia cards of the Kepler architecture (16). The Intel Xeon Phi co-processor plays just a minor role and AMD is not in use in HPC.

As expected, the only operating system running on all clusters, is Linux. The most used distribution is CentOS (Version 7 (19)), which is followed by RHEL (9) and Scientific Linux (4). It is interesting to see, that the RedHat distribution and derivatives strongly dominate the national HPC market.

Slurm is the batch system, which is installed in most cases (16 instances), followed by TORQUE (11) and LSF (3). Since all systems allow gathering system metrics like I/O or memory values, the method of gathering data and also the amount of gathered data (e.g., TORQUE is more verbose than Slurm) differs strongly. Another issue is, that Slurm has no possibility to collect data with prologue and epilogue scripts while this is possible in TORQUE.

The GCC and Intel compiler families dominate the compiler suites on the clusters (34 and 32), followed by the PGI compiler family (19). The Intel compiler family is often shipped and installed as a part of the Intel Parallel Studio XE. In the Cluster Edition this additionally contains Intel MPI, with which it is possible to generate a report about the network traffic. Intel MPI alone is installed on 27 clusters, so the above mentioned report generation could be done on over 80 % of all clusters. The most implemented MPI implementation is Open MPI (30) followed by Intel MPI (27) and MPICH and MVAPICH (9). Especially in the Open MPI case it should be possible to use special features, since it is installed on about 90 % of the clusters.

Regarding profiling tools and scientific software installed on the different clusters, there is a wide variety of programs installed on the clusters with only little overlap.

While this section served for a quick summary and overview of the results, the following section will describe the impact of these results on the project.

Impact on the Project

The results of the survey already have a strong impact in this early stage of the development of the architecture of our toolkit. The very heterogeneous infrastructure that the participating compute centres reveal is expected to become even more grave when considering the infrastructure across all computing centres in Germany. Most of these results lead to the conclusion, that a toolkit based alone on already present tools, programs and architectures is not possible.

The result of the question regarding the existing profiling tools and programs is, that Intel VTune is very widely spread. But since it is very invasive, which is in contradiction to our goal to create a lightweight toolkit, it can not seriously be considered as a part of our toolkit. Another often stated tool is vmstat, which is free and in most cases part of the distribution. This is also the case for perf in newer Linux kernels, which are not present in all computing centres yet. Another tool, which allows to collect metrics in an interval of seconds is sar. This tool collects many metrics and creates a report at run time, which can be used in our tool. To collect I/O data (disk data) we envisaged Darshan, but this is not installed on any cluster. In addition there is no clear dominance of used file systems on which we could concentrate for the choice of the tools. As a consequence, the administrators would have to install several tools on the clusters, which again is not in accordance with our goals. Thus, we will have to envisage a very modular approach or include everything needed for the desired reports in one tool, such that only little software installation is needed. In addition, this information is very valuable for one of the most important parts of the project: the report generation. Knowing that there is a wide variety of tools installed, it will be investigated whether it is possible to integrate this into the basic report and the suggestions on how to procede with application-specific profiling.

The overview of available scientific software on the compute centres form a basis on which the programs or software included in the best practices will be chosen. Here, suggestions on how to improve the runtime behaviour will be made. The most popular candidates are Matlab (15), Gaussian, OpenFOAM (both 12), VASP, ANSYS (both 9), Gromacs (8) and Turbomole, R (both 7). It is future work to develop best practices proposals for those programs.

Survey Questions and Raw Results

The questions of the survey are all listed here within the different categories, we defined. Whenever the raw data of the answers ist not used in the main body of this document, the raw data is shown here.

Hardware

1. Name of the computing centre or name of the university or of the institute, where the cluster is hosted

2. Which CPUs are implemented on the compute nodes?

Intel:

CPU	Number of answers	CPU	Number of answers
Intel E5-2609	1	Intel E5-2698 v3	1
Intel E5-2620 v4	1	Intel E5-2699 v4	1
Intel E5-2623 v4	1	Intel E5-4617	1
Intel E5-2630 v2	1	Intel E5-4620	1
Intel E5-2630 v3	5	Intel E5-4620 v2	2
Intel E5-2630 v4	1	Intel E5-4620 v3	4
Intel E5-2637 v3	1	Intel E5-4620 v4	1
Intel E5-2640	1	Intel E5-4640	1
Intel E5-2640 v3	4	Intel E5-4650	3
Intel E5-2640 v4	1	Intel E5-4655 v4	1
Intel E5-2643 v4	1	Intel E5-4670	1
Intel E5-2650	1	Intel E7-4830	2
Intel E5-2650 v4	4	Intel E5-4890 v2	1
Intel E5-2660 v2	1	Intel E7-8837	1
Intel E5-2667 v3	1	Intel E7-8860 v4	1
Intel E5-2667 v4	1	Intel E7-8891 v4	1
Intel E5-2670	7	Intel E5450	1
Intel E5-2670 v2	4	Intel E5620	1
Intel E5-2670 v3	1	Intel E7340	1
Intel E5-2680 v3	8	Intel E7540	1
Intel E5-2680 v4	1	Intel X5550	1
Intel E5-2695 v2	1	Intel X5650	2
Intel E5-2695 v4	2	Intel X5670	3
Intel E5-2697 v4	1	Intel X5675	2

The *v2, v3, v4* notations have the following meanings:

- v2 = Intel Xeon Ivy Bridge,
- v3 = Intel Xeon Haswell,
- v4 = Intel Xeon Broadwell.

AMD:

CPU	Number of answers
AMD Opteron 248	1
AMD Opteron 2376	1
AMD Opteron 2435	1
AMD Opteron 6140	1
AMD Opteron 6276	3
AMD Opteron 6308	1
AMD Opteron 6348	1
AMD Opteron 6376	2
AMD Opteron 6378	1

3. Which kinds of cluster interconnects are in use?

Interconnect type	Number of answers
Gigabit Ethernet	14
10 Gigabit Ethernet	2
Infiniband (DDR)	1
Infiniband (QDR)	19
Infiniband (FDR-10)	2
Infiniband (FDR)	10
Infiniband (EDR)	1
Aries Interconnect	1
Intel Omni-Path	4

4. Which parallel file system is used?

5. Are accelerator cards available on specific nodes on site?

Nvidia:

Accelerator Nvidia	Architecture	Number of answers
Tesla	Tesla	1
Nvidia GTX 580	Fermi	1
Tesla M2070	Fermi	1
Tesla M2075	Fermi	1
Tesla M2090	Fermi	2
Tesla K20	Kepler	2
Tesla K20m	Kepler	1
Tesla K20X	Kepler	1
Tesla K40	Kepler	2
Tesla K40m	Kepler	1
Tesla K80	Kepler	6
Kepler	Kepler	3
Maxwell	Maxwell	1
Nvidia GTX 980	Maxwell	1
Nvidia GTX 980 Ti	Maxwell	1
Nvidia GTX 1080	Pascal	1
Pascal P100	Pascal	2
Not specified	—	2
No Nvidia accelerator	—	14

Intel Xeon Phi:

Accelerator Intel Xeon Phi	Number of answers
Knights Corner	5
Knights Landing	1
Not known	1

6. Are non-POSIX memory systems available (for example object stores)?

Memory system	Number of answers
Object storage	1
Others	1
No	32

Operating system

7. Which operating system is used, which distribution and which kernel version?

OS / Distribution	Number of answers
Linux / SLES 11	2
Linux / SLES 12.1	1
Linux / Red Hat Enterprise Linux Server 6.2	1
Linux / Red Hat Enterprise Linux Server 6.6	1
Linux / Red Hat Enterprise Linux Server 6.8	2
Linux / Red Hat Enterprise Linux Server 7	2
Linux / Red Hat Enterprise Linux Server 7.2	1
Linux / Red Hat Enterprise Linux Server 7.3	2
Linux / CentOS	1
Linux / CentOS 5.5	1
Linux / CentOS 6	1
Linux / CentOS 6.8	1
Linux / CentOS 7	7
Linux / CentOS 7.1	1
Linux / CentOS 7.2	1
Linux / CentOS 7.3	6
Linux / Scientific Linux 6.8	1
Linux / Scientific Linux 7.2	3
Linux / Ubuntu 14.04 LTS	1
No answer	1

Kernel Version	Number of answers
2.6.18	1
2.6.32	2
3.0	1
3.0.101	1
3.10	15
3.12.67	1
4.0	1
4.1	1
4.4	1

8. Does a proprietary Linux environment exist?

9. Are the institutions using the container or virtualisation technique?

10. Are cgroups used?

Batch system and Workload manager

11. Which batch system is used?

Batchsystem	Number of answers
LSF	3
NQSII	1
OpenCCS	1
PBS Pro	1
SGE 6.2u5	1
Slurm	16
TORQUE (with Maui)	7
TORQUE (with Moab)	4
No answer	0

12. Usage of the nodes of the cluster: Shared or non-shared?

C and Fortran Compilers

13. Which C or Fortran compiler families are on site?

Compilerfamilies	Number of answers
ARM	0
Cray	2
GNU Compiler Collection (GNU)	34
IBM	0
Intel	32
LLVM	6
PGI	19
Open64	1
NAG (Fortran)	2
Java	1
Oracle / Solaris Studio	2
Intel Cilk Plus	1
Intel OpenCL	1
Xilinx XSDK	1
UPC	1

Parallel Program Development

14. Is Intel Parallel Studio XE installed and in which edition?

15. Which MPI implementations are installed?

MPI implementation	Number of answers
BullX MPI	1
Cray MPI	2
HP MPI	1
IBM Spectrum MPI	2
Intel MPI	27
Mellanox HPC-X Toolkit	1
MPICH1	4
MPICH2	9
MVAPICH	9
Open MPI	30
Parastation MPI	1

Profiling and Monitoring Tools

16. Which profiling and monitoring tools are installed on the cluster?

Profiling tool	Number of answers
Allinea MAP	5
Allinea Reports	5
Clustware	3
Darshan	0
Ganglia	9
Intel VTune	19
Likwid	11
PAPI	6
perf	3
Scalasca	7
Score-P	6
Slurm Profile Accounting Plugin API	6
sysstat	9
VampirTrace	9
valgrind	8
vmstat	13
None	2

Scientific Application Software and Languages

17. Which scientific applications are installed and used on your systems?

Softwarepackage	Number of answers
ABAQUS	6
ANSYS	9
Comsol	4
CP2K	4
Espresso	3
Gaussian	12
Gromacs	8
Matlab	15
NAMD	5
OpenFOAM	12
Orca	3
Python	5
R	7
Star CCM+	3
Turbomole	7
VASP	9
None	4