BigStorage: MSCA-ITN-2014-ETN-642963

Storage-based convergence between HPC and Cloud to handle Big Data

Deliverable number	D1.1
Deliverable title	Requirement description
Main Authors	Toni Cortés, all ESRs

Grant Agreement number	642963
Project ref. no	MSCA-ITN-2014-ETN-642963
Project acronym	BigStorage
Project full name	BigStorage: Storage-based convergence between HPC and Cloud to handle Big Data
Starting date (dur.)	1/1/2015 (48 months)
Ending date	31/12/2018
Project website	http://www.bigstorage-project.eu

Coordinator	María S. Pérez
Address	Campus de Montegancedo sn. 28660 Boadilla del Monte, Madrid, Spain
Reply to	mperez@fi.upm.es
Phone	+34-91-336-7380

Document Identifier	D1.1
Class Deliverable	Document
Version	1.0
Document due date	M22
Submitted	24/10/2016
Responsible	Toni Cortes, BSC
Reply to	Toni.cortes@bsc.es
Document status	final
Nature	R(Report)
Dissemination level	(Public)
WP/Task responsible(s)	Toni Cortes - BSC
Contributors	All ESRs
Distribution List	Consortium Partners
Reviewers	All advisors
Document Location	http://bigstorage-project.eu/index.php/deliverables

EXECUTIVE SUMMARY

This document provides an overview of the progress of the work done during the first 22-month period of the Project BigStorage (from 01-01-2015 until 30-09-2016) with respect to WP1.During this period, ESRs have organized themselves into four working groups, where each of these groups analysed the requirements of the for main use cases of the project: Human Brain project, Square Kilometre Array, Climate modelling, and Smart Cities.

In this deliverable we briefly describe each of the projects and list the requirements each of these projects have in the areas of covered by the ETN: Storage, IO, analysis, etc. For each requirement we give information about the requirements, it potential evolution, and sources of information where the reader can get deeper insight of the requirement. Finally, also for each project and requirement, we list what ESRs are working on solution that may cover fully or partially the listed requirements.

DOCUMENT INFORMATION

IST Project Number	MSCA-ITN-2014-ETN-642963	Acronym	BigStorage
Full Title	BigStorage: Storage-based convergence between HPC and Cloud to handle Big Data		
Project URL	http://www.bigstorage-project.eu		
Document URL	http://bigstorage-project.eu/index.php/deliverables		
EU Project Officer	Szymon Sroda		

Deliverable	Number	D1.1	Title	Requirement description
Workpackage	Number	WP1	Title	Use Case Analysis and Evaluation

Date of Delivery	Contractual	M22	Actual	24/10/2016
Status	version 3		final ■	
Nature	prototype □ report ■ dissemination □			
Dissemination level	public ■ consortium □			

Authors (Partner)	<name and institution>			
Responsible Author	Name	Toni Cortes	E-mail	Toni.cortes@bsc.es
	Partner	BSC	Phone	+34 934137966

Abstract (for dissemination)	In this deliverable we briefly describe each of the projects and list the requirements each of these projects have in the areas of covered by the ETN: Storage, IO, analysis, etc. For each requirement we give information about the requirements, it potential evolution, and sources of information where the reader can get deeper insight of the requirement. Finally, also for each project and requirement, we list what ESRs are working on solution that may cover fully or partially the listed requirements.
Keywords	Requirements, Use cases, HBP, SKA, Climate Modelling, Smart Cities, Storage, IO, Analyisis.

Version	Modification(s)	Date	Author(s)
01	The 4 WG have added the initial contents	<15/09/2016>	All ESRs
02	Put the document together, homogenization, and first review	<20/09/2016>	Toni Cortes
03	Added comments by Reviewers	<28/09/2016>	Toni Cortes

PROJECT CONSORTIUM INFORMATION

Participants	Contact
Universidad Politécnica de Madrid (UPM), Spain	María S. Pérez Email: mperez@fi.upm.es
Barcelona Supercomputing Center (BSC), Spain	Toni Cortes Email: toni.cortes@bsc.es
Johannes Gutenberg University (JGU) Mainz, Germany	André Brinkmann Email: brinkman@uni-mainz.de
Inria, France	Gabriel Antoniu Email: gabriel.antoniu@inria.fr Adrian Lebre Email: adrien.lebre@inria.fr
Foundation for Research and Technology - Hellas (FORTH), Greece	Angelos Bilas Email: bilas@ics.forth.gr

Seagate, UK		Malcolm Muggeridge Email: malcolm.muggeridge@seagate.com
DKRZ, Germany		Thomas Ludwig Email: ludwig@dkrz.de
CA Technologies Development Spain (CA), Spain		Victor Muntes Email: Victor.Muntes@ca.com
CEA, France		Jacque Charles Lafoucriere Email: Charles.LAFOUCRIERE@CEA.FR
Fujitsu Technology Solutions GMBH, Germany		Sepp Stieger Email: sepp.stieger@ts.fujitsu.com

TABLE OF CONTENTS

INTRODUCTION

One of the main tasks of a researcher is first to understand, in enough detail, the problem that he or she is trying to solve in order to avoid inventing problems to later solve them. In order to understand some of the problems that current research is facing, ESRs have analysed four key projects in the research agenda in Europe: Human Brain Project, Square Kilometre Array, Climate modelling, and Smart cities. All four projects have challenging requirements in both storage and analysis technologies, the main topics of this ETN.

In order to understand the requirements, the ESRs have read all available documentation (which is quite heterogeneous) and have discussed the topics with some relevant researchers in each community. In addition, and in order to ease the task of understanding the needs of each project, some key researchers were invited to deliver talks in the BigStorage initial school held in Barcelona.

This analysis work has been performed in a distributed way. We have built 4 working groups composed by 3 to 5 ESRs and 1 to 3 advisors. Each of these working groups have been assigned to one project and delivered the list of requirements. This task has reduced the effort of requirement gathering, and has helped ESRs to work in groups with distributed tasks.

The result of this work is a list of requirements (with a number of references to enable further details that do not fit in this deliverable, but that may be important for ESRs while performing their research). It is important to understand that, given the heterogeneity of the sources and the distributed manner of the research work, there is not a perfect matching in the level of details between the different projects. Nevertheless, we believe it is a great tool for ESRs to understand current problems and how they have been focused by the researchers in each of the projects.

In addition to the requirement list, ESRs have also tried to identify which of these described requirements may be affected by the work they are performing. This table will be the seed for the follow up work, where some of the listed requirements will be converted into benchmarks to test the progresses of the research done by ESRs with respect to the four analysed projects.

This deliverable is organized as follows: Section 2, presents the different use cases, requirements, and potential implication of the different ESRs in each requirement and Section 3 Concludes the analysis.

Use cases

Human Brain Project

Description

The Human Brain Project (HBP) is a European Commission Future and Emerging Technologies Flagship Project. It aims to put in place a cutting-edge, ICT-based scientific research infrastructure, which will allow scientific and industrial researchers to advance knowledge in the fields of neuroscience, computing and brain-related medicine. The Project promotes collaboration across the globe and is committed to drive forward European industry.

The Neuroscience Subprojects will extend their research in brain organisation and theory, in order to support the building of increasingly sophisticated models and simulations. In addition, related work will be done in brain-like computing and robotics, working up to replication of the whole mouse brain, while also laying the foundations for simulation of the much larger and more complex human brain. The resulting knowledge on the structure and connectivity of the brain will open up new perspectives for the development of "neuromorphic" computing systems incorporating unique characteristics of the brain such as energy-efficiency, fault-tolerance and the ability to learn.

The *HPC Platform Subproject* (SP7) is one of the HBP 12 operational subprojects. Its mission is to build and manage the hardware and software for the supercomputing and data infrastructure required to run cellular brain model simulations, up to the size of a full human brain. SP7 will make this infrastructure available to the consortium and the scientific community worldwide. Full brain simulations are expected to require exascale capabilities, which according to most potential suppliers' roadmaps, are likely to be available in approximately 2021-22. As well as providing sufficient computing performance, the HBP supercomputer will also need to support data-intensive interactive supercomputing and large-memory footprints. This includes topics like tightly integrated visualization, analytics, simulation capabilities, efficient European-wide data management, dynamic resource management providing co-scheduling of heterogeneous resources and a significant enlargement of memory capacity, based on power-efficient memory technologies.

Analysis/computation requirements

The main computational aim for the future is to develop ICT tools to generate high-fidelity digital reconstructions and simulations of the mouse brain and ultimately the human brain ("Simulate the Brain"). Other aims of the project include:

- Develop hardware architectures and software systems for visually interactive, multi-scale supercomputing moving towards the exascale ("Develop Interactive Supercomputing").
- Develop and operate six specialized platforms dedicated respectively to neuroinformatics, brain simulation, high performance computing, medical informatics, neuromorphic computing, neurorobotics, and a unified portal providing a single point of access to the platforms ("Develop and Operate six ICT Platforms, Making HBP Tools, Methods and Data Available to the Scientific Community").
- Develop ICT tools supporting the re-implementation of bottom-up and top-down models of the brain in neuromorphic computing and neurorobotic systems. Machine learning algorithms could also be developed, based on the HB spatial data acquired ("Develop Brain-Inspired Computing and Robotics".)
- Develop ICT tools to federate and cluster anonymized patient data ("Map Brain Diseases".)
- Implement a program of trans-disciplinary education to train young scientists to exploit the convergence between ICT and neuroscience, and to create new capabilities for European academia and industry ("Education and Knowledge Management".)

CHALLENGES

- Computation Power-

The IBM roadmap predicts the production of an exascale computer around 2018 (10^{18} flops/s). Extrapolating today's Blue Brain Project numbers, exascale is probably the minimum required to simulate the entire brain. This level of performance is just sufficient for the simultaneous computation of the present estimate of the number of equations needed to provide a first holistic version of a brain model, one that instantiates the nonlinear interactions, that give rise to the emergent properties of living brains. Regarding data storage, this is a practical problem that has effectively been solved by cloud computing and distributed storage with appropriate addressing; it is data analysis and aggregation with efficient database queries that are challenges at this scale.

- Data Gathering-

Clinical scientists are used to dealing with highly controlled, "clean" data sets, despite the messy nature of their observational constructs. Hence their data sets are often small, precious and closely guarded, being a critical part of the discovery process. This mind set is invalidated by advances in data mining algorithms that have become commonplace in industry. Such algorithms identify patterns in big data that are characterized by invariable clusters of (mathematical) rules. These powerful and computer-sensitive, data-hungry algorithms often use novel mathematics. They deal with multivariate and "dirty" data, missing data, textual or semantic data and data from different sources or with different ranges.

There are many basic science laboratory databases, often publically funded, held in universities and research laboratories around the world. These data have often been used once and exist for archival reasons alone. This mass of legacy data represents an enormous, untapped research resource. How can such heterogeneous data be usefully exploited?

Following the CERN model, asking for scientists' data in return for giving them access to many other databases should be a huge incentive, especially since it will accelerate the process of scientific discovery by increasing the efficiency of data usage.

-Computing Infrastructure-

The proposed infrastructure for the HBP can be seen in Figure 1 and some details of the sites is presented in Table 1. This kind of Infrastructure should be taken into account when proposing solutions that that the HBP requirements.

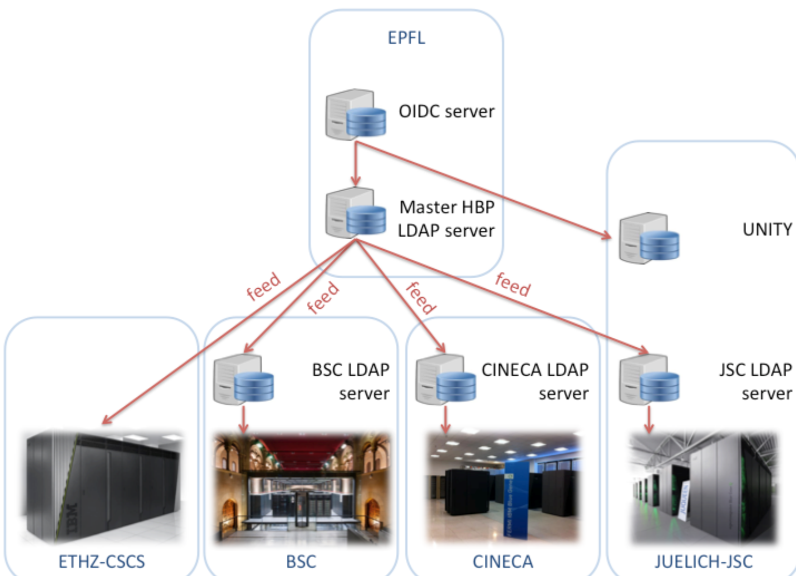

Figure 1: Infrastructure diagram used by the HBP

Name	Location	Type	Nodes	Cores	Petaflops
Juqueen	Germany	IBM/BlueGene/Q	28,672	458,752	5.9
Jureca	Germany	T-Platforms V-class architecture	1872	45,216	2.44

Piz Daint	Switzerland	Cray XC30	5,272	42,176	7.787
MareNostrum III	Spain	IBM iDataPlex	3,056	387.136	1.1
Fermi	Italy	BM/BlueGene/Q	10,240	163,840	2.1
Pico	Italy	Custom	74	1080	

Table 1: Site information for the HBP Infrastrcuture

- Data management Platforms-

- T-Storm -

T-Storm is a platform for supporting scalable real-time analytics of massive sets of voluminous time-series. The platform is constructed over the Apache Storm parallel dataflow engine and supports both vertical scalability (fully utilizing high-end servers and multi-core systems) and horizontal scalability (scaling across a cluster of physical machines or even incorporating virtual cloud resources).

- MonetDB -

MonetDB pioneered column-store solutions for high-performance data warehouses for business intelligence and eScience since 1993. It achieves its goal by innovations at all layers of a DBMS, e.g. a storage model based on vertical fragmentation, modern CPU-tuned query execution architecture, automatic and adaptive indices, run-time query optimization and a modular software.

REQUIREMENTS

Data Storage, Security and Management

Identifier	HBP1
Title	Spatial Data Management Technique
Description	Data generated by the brain simulator is really huge and complex in terms of analysis as every neuron has different characteristics and connections. To analyse the data on subcellular level, neuroscientists require more efficient models

	to fetch data. For a similar problem, a technique of prefetching spatial data with high accuracy based on previous queries is introduced. They used three approaches, the first being *SCOUT* for efficient prefetching, *FLAT* for efficient range query and *TOUCH* for efficient and scalable in-memory joins.
Prevision of change in time	Prefetching more initial spatial data, eventually speeding up the whole process.
References	[Stougiannis 2013] [Kozloski 2008]

Identifier	HBP2
Title	Spatial Join based on non-uniform dataset densities
Description	Spatial joins of two datasets with different densities according to their data distribution in an efficient manner can improve the performance of the application such as location of synapses in the human brain. Current strategy like TRANSFORMERS use an adaptable strategy for the data distribution. In indexing, it partitions the data, organizes it and computes connectivity by using self-spatial join. In Join, it uses adaptive walk by randomly selecting the role of each dataset as guide and follower. It then swaps the role of the dataset in runtime based on the density of dataset. It is an efficient scheme for datasets with non-uniform density and outperforms other schemes that only focuses on the dataset densities. It performs better due to its adaptive nature.
Prevision of change in time	The skewness of the datasets are characteristics of their application. This is the first approach that utilizes the concept of adaptability for two datasets. In Human Brain Project where

	each spike invokes tens of thousands synapsis, the dataset would become more dense and better adaptable techniques can be used.
References	[Pavlovic 2016] [Markram 2011]

Identifier	HBP3
Title	Multi-format ingestion
Description	HBP heavily relies on external data. Data that differ in physical location, storage medium and logical format. These data must be sanitized, unified and be federated under a common, simplified platform.
Prevision of change in time	Data formats could vary. The management of object, file or block storage is a challenge for the future.
References	[Frackowiak 2016]

Identifier	HBP4
Title	Anonymizing
Description	HBP's data is partially derived from clinical trials. To ensure patient's privacy the data must be sanitized (wipe out private information) and unified before ingested into the platform. However, this must happen without tampering the original data given that they are property of the issuing clinics.
Prevision of change in time	Privacy of patient's records is a very sensitive matter, which will be of more importance in the near future. The security and

	privacy of each patient will be an important social matter.
References	[Frackowiak 2016]

Identifier	HBP5
Title	Sharing
Description	Compared to other use cases, HBP is unique in terms of data collection. While others rely on in-house data, the HBP exploits third-party data, gathered over the last decades. To make the collaboration viable, each partner must not only share, but also be receive data. That makes HBP not only a processing platform, but a sharing platform as well. The platforms should be open-source and share not just data, but methods, results, and associated publications in order to accelerate the dissemination of both methods and results, and allow scientists to access the latest tools, insights, and results.
Prevision of change in time	Volume of data to be shared will increase, making the Cloud a possible solution for collaboration
References	[Frackowiak 2016]

Identifier	HBP6
Title	Performance impact of Global View Resilience(GVR) with integrated Non-volatile memories
Description	Checkpointing is a standard technique to overcome failures and an increase in checkpointing time limits a clusters performance. The bottleneck can be nearly removed by NVRAM as it provides

	high I/O bandwidth. Global View Resilience (GVR) introduces more classes of errors. The recovery test of multi versioning shows an efficient and flexible rollback performance on HPC system i.e. Blue Gene Active Storage (BGAS) with integrated NVM. This technique provides an efficient approach compared to other conventional checkpointing techniques and improves the HPC system i.e. Blue Gene for brain simulation.
Prevision of change in time	The performance can be further improved by utilizing better hybrid mapping schemes to better tier links between storage and compute nodes.
References	[Dun] [Vetter 2015]

Identifier	HBP7
Title	Storage Class Memory (SCM) in BlueGene Supercomputer to accelerate IO
Description	Modelling, simulating and analysing data generated by large scale complex brain models requiress high I/O performance and this needs sometime big caches as well as improvements of the underlying SSD arrays. The use of Direct Storage Access (DSA) supporting GPFS utilizes maximum bandwidth between an application and storage nodes and shows better scalability. Storage Class Memory (SCM) in IBM Blue Gene Supercomputer can further improve read/write performance and scalability.
Prevision of change in time	As the volume of data will increase with the computational environment, the need for better performance from SSDs and

	from the application layer is required.
References	[Schürmann 2014] [Strande 2012]

Identifier	HBP8
Title	Evaluating the performance and scalability of the Ceph distributed storage system
Description	The volume of data generated from HBP will increase, making the management and storage of this data a very important topic for the future. The usage of (Cloud-based) commodity hardware, along with CEPH, could help researchers perform experiments, without the utilisation of supercomputers. The usage of CEPH, a scalable software defined storage system, for the needs of the HBP, is a matter which should be examined.
Prevision of change in time	The storage of this data (for further use) is important, as the data could be used more than once for experiments and analysis, making the preservation of this data a critical matter, in order to avoid the re-mapping of the brain and possible re-calculations that might occur.
References	[Gudu 2014] [Weil 2006]

Data Analysis Requirements

Identifier	HBP9
Title	Multi aspect queries
Description	Brain's activity is affected by a large number of parameters, each of which has both temporal and spatial impact. Hence, a

	multi-dimension query engine is essential to boost the analysis. This engine must satisfy criteria like huge volume processing, structured and unstructured data, along with an easily understandable interface given that the majority of the involved scientists are not (and should not be) data scientists.
Prevision of change in time	Research on the human brain could bring different requirements, changing the way the query engine and interactive analytics work.
References	[Alagiannis 2012]

Identifier	HBP10
Title	Machine Learning (ML) Models based on the Human Brain
Description	Machine learning for the neuromorphic computing systems, could be based on the results from the analysis of the Human Brain. The algorithm that will be derived by the analysis of the human brain could be adjusted and fitted on a machine-learning algorithm and eventually on a robot.
Prevision of change in time	Many Robots will be based on machine learning techniques in the near future, in order to adapt to each situation. The usage of the results which will arise from the HBP analysis could help the development of better algorithms for ML. As robots reach their physical capabilities, the development of algorithms for cognitive robotics will be of vital importance for the further development of these machines.
References	[Carbajal 2015][Mahadevan 1996] [Jung 2007]

Linking Requirements to ESRs

	HBP1	HBP2	HBP3	HBP4	HBP5	HBP6	HBP7	HBP8	HBP9	HBP10
ESR1								X	X	
ESR2										X
ESR3										
ESR4	X	X				X				
ESR5										
ESR6										
ESR7										
ESR8			X					X		X
ESR9			X	X	X				X	
ESR10			X							X
ESR11	X	X				X	X			
ESR12						X	X	X		
ESR13							X	X		
ESR14			X		X			X		
ESR15							X	X		

Further reading

More general information on this project can be found in the following documents: [Golland 2014] [Tauheed 2013] [Lippert 2013] [Brömmel 2014] [Huerta 1993] .

The Square Kilometre Array

Description

The Square Kilometre array is an international scientific mega-project aiming to build and operate the world's largest radio-telescope. SKA antennas will be deployed in two locations in South Africa and Australia. With a total receiving area of over a square kilometre, the goal is to surpass the capabilities of current instruments by an order of magnitude or more [SKA]. The continuous stream of data that will be generated by the instrument will be processed and stored by specially designed compute centres located near the antennas' locations.

The Science Data Processor

The Science Data Processor is one of the 10 main work packages of the SKA project. The SDP encapsulates all the tasks required to design, provision and implement the necessary computing hardware, software packages and algorithms required to process the observation data into science products ready to be used by scientists. Long term storage as well as efficient distribution of these science products are also the responsibility of the SDP.

The major challenge of SDP is handling the vast amount of data that will be generated continuously. While hardware advances are expected to catch up by the time the telescope goes live, software performance is currently at 1/1000th of the required scaling. Therefore, significant research and innovation are required in order to design a true exascale capable system.

Understanding the type of data to be processed by the SDP is important in order to design a suitable storage architecture. SKA data is highly noisy, large in volumes and comes from multiple observations that contain incomplete samples of the target visibility. This means that the same target will be processed through multiple iterations, creating the need for temporary storage of intermediate science products. On the other hand, the data is inherently parallelizable, allowing independent processing of partial data. This leads to the idea of "compute islands" which will be formed by partially independent compute clusters responsible for processing a subset of the incoming data stream.

It is almost certain that a multi-tiered approach will be required, combining high and low performance storage elements. Based on the expected data and computation characteristics we identify four main functions that the storage stack of the SDP will have to fulfil.

1. Temporary Storage of Raw Data
2. Temporary Storage of Intermediate Data Products
3. Temporary Persistence of key data/program state

4. Archiving of science data

Based on the above we express a set of requirements for managing the efficient function of the multiple storage tiers. Maximizing efficiency of data movements and optimizing for every function is the main objective.

REQUIREMENTS

Data management

Identifier	SKA1
Title	Temporary storage of Raw Data
Description	The SDP will ingest raw data at a rate of 1 terabyte / s which requires storing the raw data temporarily to perform near-real-time processing.
Prevision	No changes
References	[Lee 1996]

Identifier	SKA2
Title	Temporary Storage of Intermediate Data Products
Description	Data products will be in this tier will be written once and read many times, either for further processing or to move to lower tiers for permanent storage.
Prevision	No changes
References	[SPD 2013][SPD 2015]

Identifier	SKA3

Title	Archiving of Science Data
Description	The challenge of the archiving tier is the very long expected lifetime of the project (>50 years).
Prevision	No changes
References	[SPD 2013][SPD 2015]

Identifier	SKA4
Title	Hierarchical Storage Management
Description	Hierarchical Storage Management solutions are currently focused mainly on archiving. A multi-tiered solution is required to efficiently manage data movement between high and low performance storage hardware.
Prevision	Tiers are expected to change as new technologies become available
References	[SPD 2013][SPD 2015]

Data placement

Identifier	SKA5
Title	Storing raw data in the processing nodes (Data Locality)
Description	Read and write patterns as well as locality are fairly predictable. It is therefore possible to move the storage elements very close to the actual processing, greatly reducing costly data movements.

Prevision	No changes
References	[SPD 2013][SPD 2015]

Fault tolerance

Identifier	SKA6
Title	Temporary Storage of Key Data/Program State
Description	• This data will be used for failure recovery hence it will exhibit a write often / read once pattern • Data will be duplicated via a double buffer system that allows for persistence.
Prevision	No changes
References	[SPD 2013][SPD 2015]

Identifier	SKA7
Title	Durability of Data
Description	Replication should be avoided at this point. An alternative could be RDDs like abstractions.
Prevision	No changes
References	[Zaharia 2012]

Performance optimizations

Identifier	SKA8
Title	Read Latency Management
Description	Explicit prefetching needed. The prefetching specification can be either provided manually or generated automatically using static code analysis.
Prevision	No changes
References	[Ibrahim 2006][Han 2005]

Identifier	SKA9
Title	Specific data structures for massively parallel accesses
Description	Today's data-structures do not fully exploit multi-core machines parallelism --> identify suitable solutions e.g., lock-free stores for high-read throughput
Prevision	No changes
References	[Fan 2013]

Linking Requirements to ESRs

	SKA1	SKA2	SKA3	SKA4	SKA5	SKA6	SKA7	SKA8	SKA9
ESR1	X				X				X
ESR2			X				X		
ESR3									
ESR4									
ESR5								X	
ESR6				X					
ESR7									
ESR8									
ESR9	X	X		X		X			X
ESR10									
ESR11					X				
ESR12					X				
ESR13							X		X
ESR14								X	
ESR15				X					

CLIMATE SCIENCE

DESCRIPTION

Climate Science is the study of climate and is part of the Earth system sciences. Related topics include weather forecasting, disaster prediction and climate model research. It considers and studies the origin, development and changes of weather patterns over a long period of time. This study aims to investigate past climate, forecast future climate conditions and to see what influence it has on various aspects of our life on Earth.

Climate change has a known impact on different components of Earth system and substantial effect on human health in particular [CCS]. It drastically affects on water resources, ecosystems, food production and agriculture, industry, settlement and society, urban life, energy consumption and many other important things of our habitat and planet in whole [SGC]. The availability of information about weather and climate conditions (temperature, winds, rains, humidity, snow, thunderstorms etc.) on a daily basis became very significant for people in analysis, planning and decision making.

To predict future changes and their consequences, earth system scientists develop and maintain climate models with high-end simulations for long periods of time. These useful research tools help to better understand present and past observed climates, support experiments under given boundary conditions of anthropogenic and natural forces. Unlike these models, similar numerical weather prediction models are focused on having current meteo-information about atmospheric conditions and forecasting the future state of weather on a day-to-day basis through simulations for short periods of time.

Standard Earth system models have a comprehensive structure that consists of many components which are responsible for simulation and treatment of major parts in Earth's system. Among them usually are models of physical components (atmosphere, oceans, land-surface) and biogeochemical subsystems (aerosols, atmospheric chemistry, vegetation) [Heavens 2013]. Execution of such complex applications with numerical models of the climate system demands powerful computing infrastructure that will support it. Thus, climate science applications are key users of HPC which still continue to improve and drive scientific progress.

Unfortunately, the progress is limited by computer capabilities. Even though today's computing power of supercomputers is very advanced and still increases, the complexity of climate models also increases and requires more computer resources than before or even more than affordable. For example, large scale experiments with help of climate models require longer simulation runs to forecast climate change during a century which means using more processor hours. Among other main problems in HPC while working with climate models is handling of huge amounts of data produced by simulations. Climate models with high resolution would not produce valuable results as they are constrained by available resources: storage resources, proper data management and post-processing.

Currently, more and more climate models are being developed and improved by earth system scientists and computational experts together with large computing centres. It is important to work closely with climate scientists to improve scientific applications and find the best computational methods and algorithms, programming languages and techniques. All this will foster the improvement of current supercomputer systems and enhance the development of Earth system models

There are a number of different centres and research institutes dedicated to climate research. Among them are German Climate Computing Center (DKRZ) and Max Planck Institute for Meteorology (MPI-M) in Hamburg (Germany), European Centre for Medium-Range Weather Forecasts (ECMWF) in Reading (United Kingdom), Barcelona Supercomputing Center (BSC) in Barcelona (Spain) which investigate questions related to climate change, physical processes in the atmosphere and perform a significant work in climate modelling. Their requirements are different due to the variety of analysis techniques and objectives. For example, ECMWF uses a significant historical archival to replay and update the models, while having tight deadlines to process forecasts. In the next section, a list of requirements regarding the computing infrastructure captured from various models and centres can be found. They were discussed together with scientist from the climate modelling community and are outlined here as the most important ones towards obtaining meaningful, credible and useful scientific results.

We identified multiple enabling requirements for climate science data storage and processing, which we divide into seven categories:

- **I/O requirements** enabling high-performance, parallel access to the data [CLM1, CLM2]
- **Data management requirements** necessary to ensure that both the large volumes of source and result data remain accessible indefinitely for reprocessing or further processing [CLM3, CLM5]
- **Data processing requirements** necessary to ensure the need for high spatial and temporal resolution in order to improve the fidelity of climate models or the predicted new era of climate forecasting [CLM4]
- **Data standards and interoperability requirements** guaranteeing that the stored data remains usable and easily linkable to new experiments if necessary, thanks to well-defined data/metadata standards and data quality assessment [CLM6, CLM7]
- **Climate data acquisition to data access requirements** to ensure that meaningful data can be used and shared efficiently amongst scientists and institutions around the world [CLM8]

- **Software infrastructure requirements** allowing the source code used to perform experiments to remain accessible, usable, and maintainable and over time [CLM9]
- **Social requirements** to foster trans-domain collaboration between climate scientists and software engineers [CLM10]

REQUIREMENTS

I/O requirements

Identifier	CLM1
Title	Parallel I/O
Description	Serial reading and writing enormous amount of data for most of today's climate science applications is a serious bottleneck during simulation performance on supercomputers. Sequential I/O usage has a great affect on their overall performance. Adopting parallel I/O is the solution to reduce or remove this bottleneck with computational costs. Parallelization in this case provides simultaneous access for each process in application either to one file or to many files.
Prevision	I/O requirements increase (both for throughput and latency)
References	[Dennis 2011][Kuhn 2013] [Henderson 1994] [Fu 2010]

Identifier	CLM2
Title	Parallel distributed file systems exploitation
Description	Parallel I/O is an essential requirement to leverage the heavily parallel computation capabilities of modern computing platforms, both on HPC platforms and modern clouds. In contrast, sequential I/O would lack the required scalability for most climate applications today on these platforms.
Prevision	Exploitation of parallel DFS is inevitable in order to use efficiently parallel I/O

References	[Dennis 2011] [Huang 2014] [Kuhn 2013]

Data management requirements

Identifier	CLM3
Title	Long term archiving of climate data
Description	Collecting as much data as possible is key to improve the likelihood of achieving a scientific discovery of climate change. Long-term archiving entails easy, affordable and timely access for a large number of scientists spanning many different fields, while additionally enabling reprocessing large data sets as understanding of sensor performance, algorithms and earth science improves. Examples of new information that would warrant long-term archival for data reprocessing are detection of sensor calibration drift, the availability of ancillary data sets, better climate models, or simply errors in previous processing. Additionally, the data must to stay readable indefinitely, specifically by ensuring that it does not get corrupted or lost over time.
Prevision	Volume (number of archives) will increase
References	[Luthardt 2015] [Cash 2015]

Identifier	CLM4
Title	Exabyte-scale data storage capacity
Description	In climatology, there are huge volumes of data from observation and climate model simulations. These data are critical for understanding the climate mechanisms and predicting the future climate change. Such data intensive experiments are expected to generate exabytes of data over the next 5-10 years, which must be transferred, visualized, and analysed by geographically distributed teams of researchers. The amount of data collected and produced is expanding at a staggering rate, and projected to exceed hundreds of exabytes by

	2020.
Prevision	Data volume will inevitably increase with a high speed and climate applications will demand more storage space for it.
References	[Dart 2007] [Fernández-Quiruelas 2011] [André 2012]

Identifier	CLM5
Title	Capacity over performance
Description	Such archival is necessary for producing high-quality output and fosters scientific advances in the field of climatology. Using cost-effective, highly-reliable and resilient storage systems is mandatory to achieve such long-term archival using a wide variety of mediums such as spinning disks or magnetic tapes. In comparison, high-performance storage solutions based on flash storage trade both price, volume and resilience for performance. As data access times is less important than the available volume of data for long-term archives, high-capacity archives must be used, and systems must prioritize high capacities over high performance.
Prevision	Storage capacity will be in high demand
References	[Luthardt 2015] [Jewell 2014]

Data standards and Interoperability

Identifier	CLM6
Title	Data and metadata standardization
Description	In order to enable easily linking and comparing current data with archived data when running new experiments, the data must be structured with well-defined

	standards which are used in climate science and more generally current HPC applications. Examples of such standards are NetCDF and HDF. Such standardization must be applied to both data and metadata.
Prevision	All produced data and metadata sets will fall under general standard otherwise different formats will appear
References	[CCDF] [FDSD] [Luthardt 2015]

Identifier	CLM7
Title	Data quality assessment and curation
Description	Data storage is confined to simply keeping data in existence and ensuring that it can be accessed when needed. As such, data curation is also necessary to entail practices of refreshment or format migration (essential to maintaining the data in a usable form for re-running experiments or linking new ones with archived data) and to call for higher-level curatorial practices such as enhancement of the data through added metadata, or migration from one representational standard to another.
Prevision	Strengthened data quality requirements
References	[Luthardt 2015]

Climate Data access requirements

Identifier	CLM8
Title	Scientific data accessibility and dissemination
Description	Produced and collected volumes of climate data are very useful and important for the scientific community in further research projects. Restricted access to

	data which have not been evaluated through quality control procedure will only reduce the velocity of progress and benefits of climate change research work. Thus, it is required that valuable climate datasets must have digital object identifier (DOI), should be stored in and disseminated through domain-specific repositories where researchers can download them via the Internet free of charge. Every repository must be assessed and approved to ensure digital data producers that their published information is safe and properly managed. Although some repositories might charge for scientific data hosting, all materials have to be freely available for search and retrieving in non-commercial, research and educational purposes. However, users who want to exploit climate datasets in research work might be asked to apply (optionally and without any costs) for an access permission. It will help to validate the purposes of data usage and to see statistics on how they serve the scientific community. Users (for example, researchers) in their turn are obliged to make proper data citation and acknowledgements in their papers, articles etc.
Prevision	Scientific data restrictions elimination
References	[Luthardt 2015] [RDR] [DSA] [IS-ENESD6.1] [Overpeck 2011]

Software infrastructure requirements

Identifier	CLM9
Title	Model code compatibility
Description	In order to cope with ever-evolving tools, libraries, architectures and hardware, the code used for scientific experiments in climate science must be archived with the data, and kept backwards compatible to guarantee that even old code will remain runnable on new machines. Consequently, the code should rely on standard tools rather than platform-specific libraries (such as CUDA). This enables data reprocessing at a later stage regardless of the platform changes that could have occurred since it has been first developed.

Prevision	Existing code should run on new hardware with relatively little effort otherwise code refactoring will be needed
References	[Intel] [Kindratenko 2009] [SED]

Social requirements

Identifier	CLM10
Title	Scientific communication and collaboration with climate community
Description	Leveraging the latest advances in computer science is not an easy task for climate scientists, as the developed solutions are rarely usable out-of-the-box for climate science, consequently leading to slow or limited adoption of HPC innovations by the climate community. Similarly, the limited knowledge of computer scientists on climate science significantly limits understanding of the requirements and constraints of climate research. In order to efficiently employ all advances in computer science for climate research, and to enable purpose-designed applications to evolve along with latest model finding from climate scientists, it is of crucial importance to create bridges between these two communities. This can take the form of trainings, active and continuous collaboration, as well as federated access to data and models across the two communities.
Prevision	Collaboration between computer and climate scientists will be necessary in order to make next generation of HPC useful for Earth science modelling
References	[Mitchell 2012] [Washington 2009]

Linking Requirements to ESRs

	CLM1	CLM2	CLM3	CLM4	CLM5	CLM6	CLM7	CLM8	CLM9	CLM10
ESR1		X		X		X		X		X
ESR2										
ESR3	X	X		X						X
ESR4		X								
ESR5										
ESR6										
ESR7		X		X	X					X
ESR8			X							
ESR9	X			X	X			X		
ESR10										
ESR11										
ESR12	X	X								
ESR13										
ESR14	X	X		X		X				X
ESR15										

SMART CITIES

DESCRIPTION

The proliferation of small sensors and devices around us capable of generating valuable information has created a new paradigm of computing known as Internet of Things. One of the most important concepts of this paradigm is the possibility of integrating and analysing all of this data input to make smart decisions in fields like healthcare, traffic management, water quality, air pollution and many more [Gubbi 2013].

Smart Cities are able to take advantage of these IoT networks to improve citizen's life. It can bring benefits to public transportation, garbage management, parking or education to name some examples [Zanella 2014]. The objective is to be able to gather data from different sensors located around the Smart City infrastructure and analyse it either at real time or offline. The information gathered can be used to improve services, take decisions or be published as open data for the citizens.

However, building this kind of infrastructure poses new challenges in many different levels. First the data has important contextual information like spatial or temporal that has to be considered at the moment of storing it. There are also privacy and security concerns that can be raised to protect the privacy of citizens. Also the speed and volume of the data generated has to be considered. But most important is the heterogeneity of the different sensors and systems involved in this kind of infrastructure. For example, the communication protocol or the architecture of the sensors/devices can be different depending on its type creating the necessity of integrating all of the different sources and storing the data in our frontend.

Measurements and data coming from IoT devices are not only processed in the cloud, since the infrastructure and processing capabilities can be insufficient. The needs of, e.g., geographical distribution of resources, real-time communication, incorporation with large networks are handled by fog computing. Through this paradigm, part of processing is done by edge devices or clouds closer to data sources, resulting in less latency and bandwidth usage [Vaquero 2014].

Lastly the huge amount of sensors and devices involved makes it difficult to detect failures and problems in the network, like disconnected wires, high and abnormal energy consumption or wrongly configured devices. The detection and root cause analysis of these events have to be considered to operate a Smart City effectively.

Internet of Things and Smart Cities is a novel area that is attracting lots of interest and creating new research challenges. Here we try to detect the requirements needed by this kind of networks and linking them to the different BigStorage objectives.

REQUIREMENTS

Analysis/computation requirements

Identifier	SCT1
Title	Stream analysis
Description	Data could be analysed in real time to monitor different aspects of the city (Environment, traffic…)
Prevision	Streams of data is closely related to the use case
References	[Kitchin 2014]

Identifier	SCT2
Title	Spatial and temporal data
Description	The nature of the data generated through sensors has embedded spatial and temporal data (e.g. When was the measure generated and where?)
Prevision	No changes
References	[Gubbi 2013]

Identifier	SCT3
Title	Open and accessible data

Description	This huge amounts of data have to be open and/or accessible for its use. This also brings privacy and security challenges
Prevision	No changes
References	[ODBDW]

Identifier	SCT4
Title	Batch processing and learning from data
Description	In addition to real-time data processing huge amounts of data can be also analysed off-line (optimising public transport routes, etc.).
Prevision	No changes expected
References	[Zanella 2014]

Storage requirements

Identifier	SCT5
Title	Storage in real time
Description	Multiple sensors generate data with high velocity that has to be stored almost in real time
Prevision	No changes expected
References	[Kitchin 2014]

Identifier	SCT6
Title	Replicated storage system
Description	Dependability vis provision of replicated storage
Prevision	No changes expected
References	[Nastic 2014]

Infrastructure requirements

Identifier	SCT7
Title	Heterogeneous environment
Description	The architecture of a Smart City involves connecting heterogeneous environments with different protocols and technologies (sensors, storage system, backend, frontend...)
Prevision	No changes expected
References	[Zanella 2014]

Identifier	SCT8
Title	Data locality
Description	It is not necessary to send all data around the world, but rather process it locally and send aggregates.

Prevision	No changes expected
References	[Vaquero 2014]

Identifier	SCT9
Title	Fault detection system for IoT system
Description	Detect wrongly configured devices, disconnected wires, explain accurately occurrences of combined faults. Detect and explain high energy consumption
Prevision	No changes expected
References	[Niggemann 2016] [Lazarova-Molnar 2016]

Identifier	SCT10
Title	Scalable system
Description	It has to be scalable (able to add new sensors and input sources), including the ability to ingest new data with a structure that is not known in advance.
Prevision	No changes expected
References	[Zanella 2014]

Linking Requirements to ESRs

	SCT1	SCT2	SCT3	SCT4	SCT5	SCT6	SCT7	SCT8	SCT9	SCT10
ESR1	X	X	X	X	X	X		X		X
ESR2	X			X			X	X		
ESR3										
ESR4				X						
ESR5										
ESR6										
ESR7										
ESR8					X					
ESR9		X	X		X	X	X			X
ESR10										
ESR11										
ESR12		X			X	X				X
ESR13					X	X				
ESR14					X					
ESR15							X		X	X

CONCLUSION

In this deliverable, we have presented a list of requirements from four key research projects in Europe: Human Brain Project, Square Kilometre Array, Climate modelling, and Smart Cities. We have also identified what ESR's research has more potential in affecting each of the requirements. This information will help ESRs to find good mechanisms to test the progress of their work in a real problem.

References

[Alagiannis 2012] Alagiannis, I., Borovica, R., Branco, M., Idreos, S., & Ailamaki, A. (2012, May). NoDB: efficient query execution on raw data files. In Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data (pp. 241-252). ACM. https://www.researchgate.net/publication/254006772_NoDB_Efficient_Query_Execution_on_Raw_Data_Files

[André 2012] Jean-Claude André, Two remarks on future developments of climate simulation, with strong impact on computing and data processing, http://www.exascale.org/bdec/sites/www.exascale.org.bdec/files/whitepapers/BDEC-Climate.Forecasting-JCA-15.01.12.pdf

[Brömmel 2014] Brömmel, D., Suarez, E., Orth, B., Graf, S., Detert, U., Pleiter, D., ... & Lippert, T. (2014). Paving the road towards pre-exascale supercomputing. In*NIC Symposium 2014* (No. FZJ-2014-01327). Jülich Supercomputing Center.

[Carbajal 2015] Carbajal, J. P., Dambre, J., Hermans, M., & Schrauwen, B. (2015). Memristor models for machine learning. Neural computation. http://www.mitpressjournals.org/doi/abs/10.1162/NECO_a_00694#.V99_3Ft9672

[Cash 2015] Ben Cash, Climate Modeling and Big Data: Current Challenges and Prospects for the Future, http://cra.org/wp-content/uploads/2015/08/Cash.pdf

[CCDF] Common climate data formats: overview https://climatedataguide.ucar.edu/climate-data-tools-and-analysis/common-climate-data-formats-overview

[CCS] Fact sheet: Climate change science - the status of climate change science today, https://unfccc.int/files/press/backgrounders/application/pdf/press_factsh_science.pdf

[Dart 2007] E. Dart and B. Tierney (editors), "BES Science Network Requirements", Report of the Basic Energy Sciences Network Requirements Workshop sponsored by Basic Energy Sciences Program Office, DOE Office of Science and the Energy Sciences Network 2007.

[Dennis 2011] Dennis, John M., et al. "An application-level parallel I/O library for Earth system models." International Journal of High Performance Computing Applications (2011): 1094342011428143.

[DSA] The Data Seal of Approval http://www.datasealofapproval.org/en/information/about/

[Dun] Nan Dun, Dirk Plieter, Aiman Fang, Nicolas, Andrew,"Multi Versioning Performance Opportunities in BGAS System for Resilience", University of Chicago - JSC, Julich, germany, "https://sites.google.com/site/uchicagolssg/lssg/research/gvr"

[Fan 2013] Bin Fan, David G. Andersen, and Michael Kaminsky. 2013. MemC3: compact and concurrent MemCache with dumber caching and smarter hashing. In Proceedings of the 10th USENIX conference on Networked Systems Design and Implementation (NSDI'13).

[FDSD] Formats for delivery of scientific/environmental data, https://geo-ide.noaa.gov/wiki/index.php?title=Formats_for_delivery_of_scientific/environmental_data

[Fernández-Quiruelas 2011] Fernández-Quiruelas, V., Fernández, J., Cofiño, A. S., Fita, L., & Gutiérrez, J. M. (2011). Benefits and requirements of grid computing for climate applications. An example with the community atmospheric model. Environmental Modelling & Software, 26(9), 1057-1069.

[Frackowiak 2016] Frackowiak, R., Ailamaki, A., & Kherif, F. (2016). Federating and Integrating What We Know About the Brain at All Scales: Computer Science Meets the Clinical Neurosciences. In Micro-, Meso-and Macro-Dynamics of the Brain (pp. 157-170). Springer International Publishing. link.springer.com/content/pdf/10.1007%2F978-3-319-28802-4_10.pdf

[Fu 2010] Fu, J., Liu, N., Sahni, O., Jansen, K. E., Shephard, M. S., and Carothers, C. D.: Scalable parallel I/O alternatives for massively parallel partitioned solver systems, in: Proceedings of the Workshop on Large-Scale Parallel Processing in conjunction with the IEEE International Parallel and Distributed Processing Symposium, 1–8, Atlanta, Georgia, USA, 2010.

[Golland 2014] Golland, P., Gallant, J., Hager, G., Pfister, H., Papadimitriou, C., Schaal, S., & Vogelstein, J. T. (2014). A New Age of Computing and the Brain: Report of the CCC Brain Workshop.

[Gubbi 2013]. Gubbi, Jayavardhana, et al. "Internet of Things (IoT): A vision, architectural elements, and future directions." *Future Generation Computer Systems*29.7 (2013): 1645-1660.

[Gudu 2014] Gudu, D., Hardt, M., & Streit, A. (2014, October). Evaluating the performance and scalability of the ceph distributed storage system. In Big Data (Big Data), 2014 IEEE International Conference on (pp. 177-182). IEEE. http://ieeexplore.ieee.org/document/7004229/

[Han 2005] W. Han, K. Whang, and Y. Moon, "A formal framework for prefetching based on the type-level access pattern in object-relational DBMSs," *IEEE Trans. Knowledge Data Eng.*, vol. 17, no. 10, pp. 1436–1448, 2005.

[Heavens 2013] Heavens, N. G., Ward, D. S. & Natalie, M. M. (2013) Studying and Projecting Climate Change with Earth System Models. Nature Education Knowledge 4(5):4

[Henderson 1994] Henderson, M., Nickless, B., & Stevens, R. (1994, May). A scalable high-performance I/O system. In Scalable High-Performance Computing Conference, 1994., Proceedings of the (pp. 79-86). IEEE.

[Huang 2014] Huang, X. M., Wang, W. C., Fu, H. H., Yang, G. W., Wang, B., & Zhang, C. (2014). A fast input/output library for high-resolution climate models. Geoscientific Model Development, 7(1), 93-103.

[Huerta 1993] Huerta, M. F., Koslow, S. H., & Leshner, A. I. (1993). The human brain project: an international resource. *Trends in neurosciences*, *16*(11), 436-438.

[Ibrahim 2006] A. Ibrahim and W. Cook, "Automatic prefetching by traversal profiling in object persistence architectures," in *Proceedings of the 20th European Conference on Object-Oriented Programming*, ser. ECOOP 2006. Springer-Verlag, 2006, pp. 50–73.

[Intel] Intel® Compiler for Linux*: Compatibility with GNU compilers, https://software.intel.com/sites/products/collateral/hpc/compilers/intel_linux_compiler_compatibility_with_gnu_compilers.pdf

[IS-ENESD6.1] IS-ENES2 deliverable (D-N: 6.1) Report on access rights for CMIP5 and CORDEX for commercial use, https://verc.enes.org/ISENES2/documents/deliverables/is-enes2_d6-1_report-on-access-rights-for-cmip5-and-cordex-for-commercial-use

[Jewell 2014] D. Jewell, R. D. Barros, S. Diederichs, L. M. Duijvestijn, M. Hammersley, A. Hazra, C. Holban, Y. Li, O. Osaigbovo, A. Plach, I. Portilla, M. Saptarshi, H. P. Seera, E. Stahl, and C. Zolotow. (2014, Jan.). Performance and capacity implications for big data. Int. Bus. Mach.. [Online]. Available: http://www.redbooks.ibm.com/redpapers/pdfs/redp5070.pdf

[Jung 2007] Jung, Y., Choi, Y., Park, H., Shin, W., & Myaeng, S. H. (2007, August). Integrating robot task scripts with a cognitive architecture for cognitive human-robot interactions. In 2007 IEEE International Conference on Information Reuse and Integration (pp. 152-157). IEEE.
http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.700.9117&rep=rep1&type=pdf

[Kindratenko 2009] GPU Clusters for High-Performance Computing http://www.ncsa.illinois.edu/People/kindr/papers/ppac09_paper.pdf

[Kitchin 2014] Kitchin, Rob. "The real-time city? Big data and smart urbanism." GeoJournal 79.1 (2014): 1-14.

[Kozloski 2008] Kozloski, James, et al. "Identifying, tabulating, and analyzing contacts between branched neuron morphologies." *IBM Journal of Research and Development* 52.1.2 (2008): 43-55.

[Kuhn 2013] A Semantics-Aware I/O Interface for High Performance Computing (Michael Kuhn), In Supercomputing, Lecture Notes in Computer Science (7905), pp. 408–421, (Editors: Julian Martin Kunkel, Thomas Ludwig, Hans Werner Meuer), Springer (Berlin, Heidelberg), ISC 2013, Leipzig, Germany, ISBN: 978-3-642-38749-4, ISSN: 0302-9743, 2013-06

[Lazarova-Molnar 2016] Lazarova-Molnar, S., Shaker, H. R., & Mohamed, N. (2016, March). Fault detection and diagnosis for smart buildings: State of the art, trends and challenges. In 2016 3rd MEC International Conference on Big Data and Smart City (ICBDSC) (pp. 1-7). IEEE.

[Lee 1996] C. Lee, C. Kesselman and S. Schwab, "Near-real-time satellite image processing: Metacomputing in CC++", *IEEE Comput. Graph. Appl.*, vol. 16, no. 4, pp. 79-84, 1996.

[Lippert 2013] Lippert, T., & Orth, B. (2013, July). Supercomputing Infrastructure for Simulations of the Human Brain. In *International Workshop on Brain-Inspired Computing* (pp. 198-212). Springer International Publishing.

[Luthardt 2015] Hans Luthardt, Hans Ramthun, Frank Toussaint, Long Term Archive Handbook, DKRZ – Department Data Management, Version 1.1, 2015 https://www.dkrz.de/daten-en/data-services/long_term_archiving/DKRZ-lta-handbook_rev_1.1.pdf?lang=en

[Mahadevan 1996] Mahadevan, S. (1996). Machine learning for robots: A comparison of different paradigms. In Proceedings of the Workshop on Towards Real Autonomy, IEEE/RSJ Internaltional Conference on Intelligent Robots and Systems (IROS'96). http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.40.5162&rep=rep1&type=pdf

[Markram 2011] H. Markram, K. Meier, S. Grillner, R. Frackowiak, S. Dehaene, A. Knoll, H. Sompolinsky, K. Verstreken, J. DeFelipe, S. Grant, and J.-P. Changeux, "Introducing the Human Brain Project," vol. 7, no. 0, 2011, fET' 11 http://www.sciencedirect.com/science/article/pii/S1877050911006806

[Mitchell 2012] Mitchell, J., Budich, R., Joussaume, S., Lawrence, B., & Marotzke, J. (2012). Infrastructure strategy for the European Earth system modelling community 2012–2022.

[Nastic 2014] Nastic, S., Sehic, S., Le, D. H., Truong, H. L., & Dustdar, S. (2014, August). Provisioning software-defined iot cloud systems. In Future Internet of Things and Cloud (FiCloud), 2014 International Conference on (pp. 288-295). IEEE.

[Niggemann 2016] Niggemann, Oliver, et al. "Data-Driven Monitoring of Cyber-Physical Systems Leveraging on Big Data and the Internet-of-Things for Diagnosis and Control."

[ODBDW] http://www.icsu.org/science-international/accord/open-data-in-a-big-data-world-short

[Overpeck 2011] Overpeck, J. T., Meehl, G. A., Bony, S., & Easterling, D. R. (2011). Climate data challenges in the 21st century. science, 331(6018), 700-702.

[Pavlovic 2016] Pavlovic, M., Heinis, T., Tauheed, F., Karras, P., & Ailamaki, A. (2016, May). TRANSFORMERS: Robust spatial joins on non-uniform data distributions. In2016 IEEE 32nd International Conference on Data Engineering (ICDE) (pp. 673-684). IEEE. https://spiral.imperial.ac.uk/bitstream/10044/1/30595/2/transformers_compressed.pdf

[RDR] Recommended Data Repositories http://www.nature.com/sdata/policies/repositories

[Schürmann 2014] Schürmann, Felix, et al. "Rebasing i/o for scientific computing: Leveraging storage class memory in an ibm bluegene/q supercomputer." *International Supercomputing Conference*. Springer International Publishing, 2014.

[SED] Software Engineering and Code Development for HPC Applications http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.126.8459&rep=rep1&type=pdf

[SGC] Statement of Guidance for Climate (other aspects -CCl) - An analysis of current and emerging capacity gaps in surface and upper air observations to support climate activities. http://www.wmo.int/pages/prog/www/OSY/SOG/SoG-Climate-CCl.doc

[SKA] http://www.skatelescope.org

[SDP 2013] Paul Alexander, Chris Broekema, Simon Ratcliffe, Rosie Bolton, Bojan Nikolic. "SDP Element Concept", SKA Project Document number SDP-PROP-DR-001-1 https://www.skatelescope.org/wp-content/uploads/2013/09/SDP-PROP-DR-001-1_ElemConc.pdf

[SDP 2015] P.C. Broekemaa, R.V. van Nieuwpoortb and H.E. Balc. "The Square Kilometre Array Science Data Processor. Preliminary compute platform design". Journal of Instrumentation, Volume 10, Issue 07, article id. C07004 (2015). http://dx.doi.org/10.1088/1748-0221/10/07/C07004

[Stougiannis 2013] Stougiannis, M. Pavlovic, et al. 2013. Data-driven neuroscience: enabling breakthroughs via innovative data management. In *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data* (SIGMOD '13). ACM

[Strande 2012] Strande, Shawn M., et al. "Gordon: design, performance, and experiences deploying and supporting a data intensive supercomputer." *Proceedings of the 1st Conference of the Extreme Science and Engineering Discovery Environment: Bridging from the eXtreme to the campus and beyond*. ACM, 2012.

[Tauheed 2013] Tauheed, F., Nobari, S., Biveinis, L., Heinis, T., & Ailamaki, A. (2013, September). Computational neuroscience breakthroughs through innovative data management. In *East European Conference on Advances in Databases and Information Systems* (pp. 14-27). Springer Berlin Heidelberg.

[Vaquero 2014]. L. M. Vaquero and L. Rodero-Merino, "Finding your Way in the Fog: Towards a Comprehensive Definition of Fog Computing," ACM SIGCOMM Computer Communication Review, vol. 44, no. 5, pp. 27- 32, October 2014.

[Vetter 2015] Vetter, Jeffrey S., and Sparsh Mittal. "Opportunities for nonvolatile memory systems in extreme-scale high-performance computing." *Computing in Science & Engineering* 17.2 (2015): 73-82.

[Washington 2009] Washington, W. M., Buja, L., & Craig, A. (2009). The computational future for climate and Earth system models: on the path to petaflop and beyond. Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences, 367(1890), 833-846.

[Weil 2006] Weil, S. A., Brandt, S. A., Miller, E. L., Long, D. D., & Maltzahn, C. (2006, November). Ceph: A scalable, high-performance distributed file system. InProceedings of the 7th symposium on Operating systems design and implementation (pp. 307-320). USENIX Association.

[Zaharia 2012] Matei Zaharia, Mosharaf Chowdhury, Tathagata Das, Ankur Dave, Justin Ma, Murphy McCauley, Michael J. Franklin, Scott Shenker, and Ion Stoica. 2012. Resilient distributed datasets: a fault-tolerant abstraction for in-memory cluster computing. (NSDI'12).

[Zanella 2014]. Zanella, Andrea, et al. "Internet of things for smart cities." *IEEE Internet of Things Journal* 1.1 (2014): 22-32.