



Status Report

Alberto Di Meglio – CERN openlab Head


13/04/2018

CERN OPENLAB'S MISSION

A Public-Private Partnership to Foster Research and Innovation

Evaluate and test state-of-the-art technologies in a challenging environment and improve them in collaboration with industry.

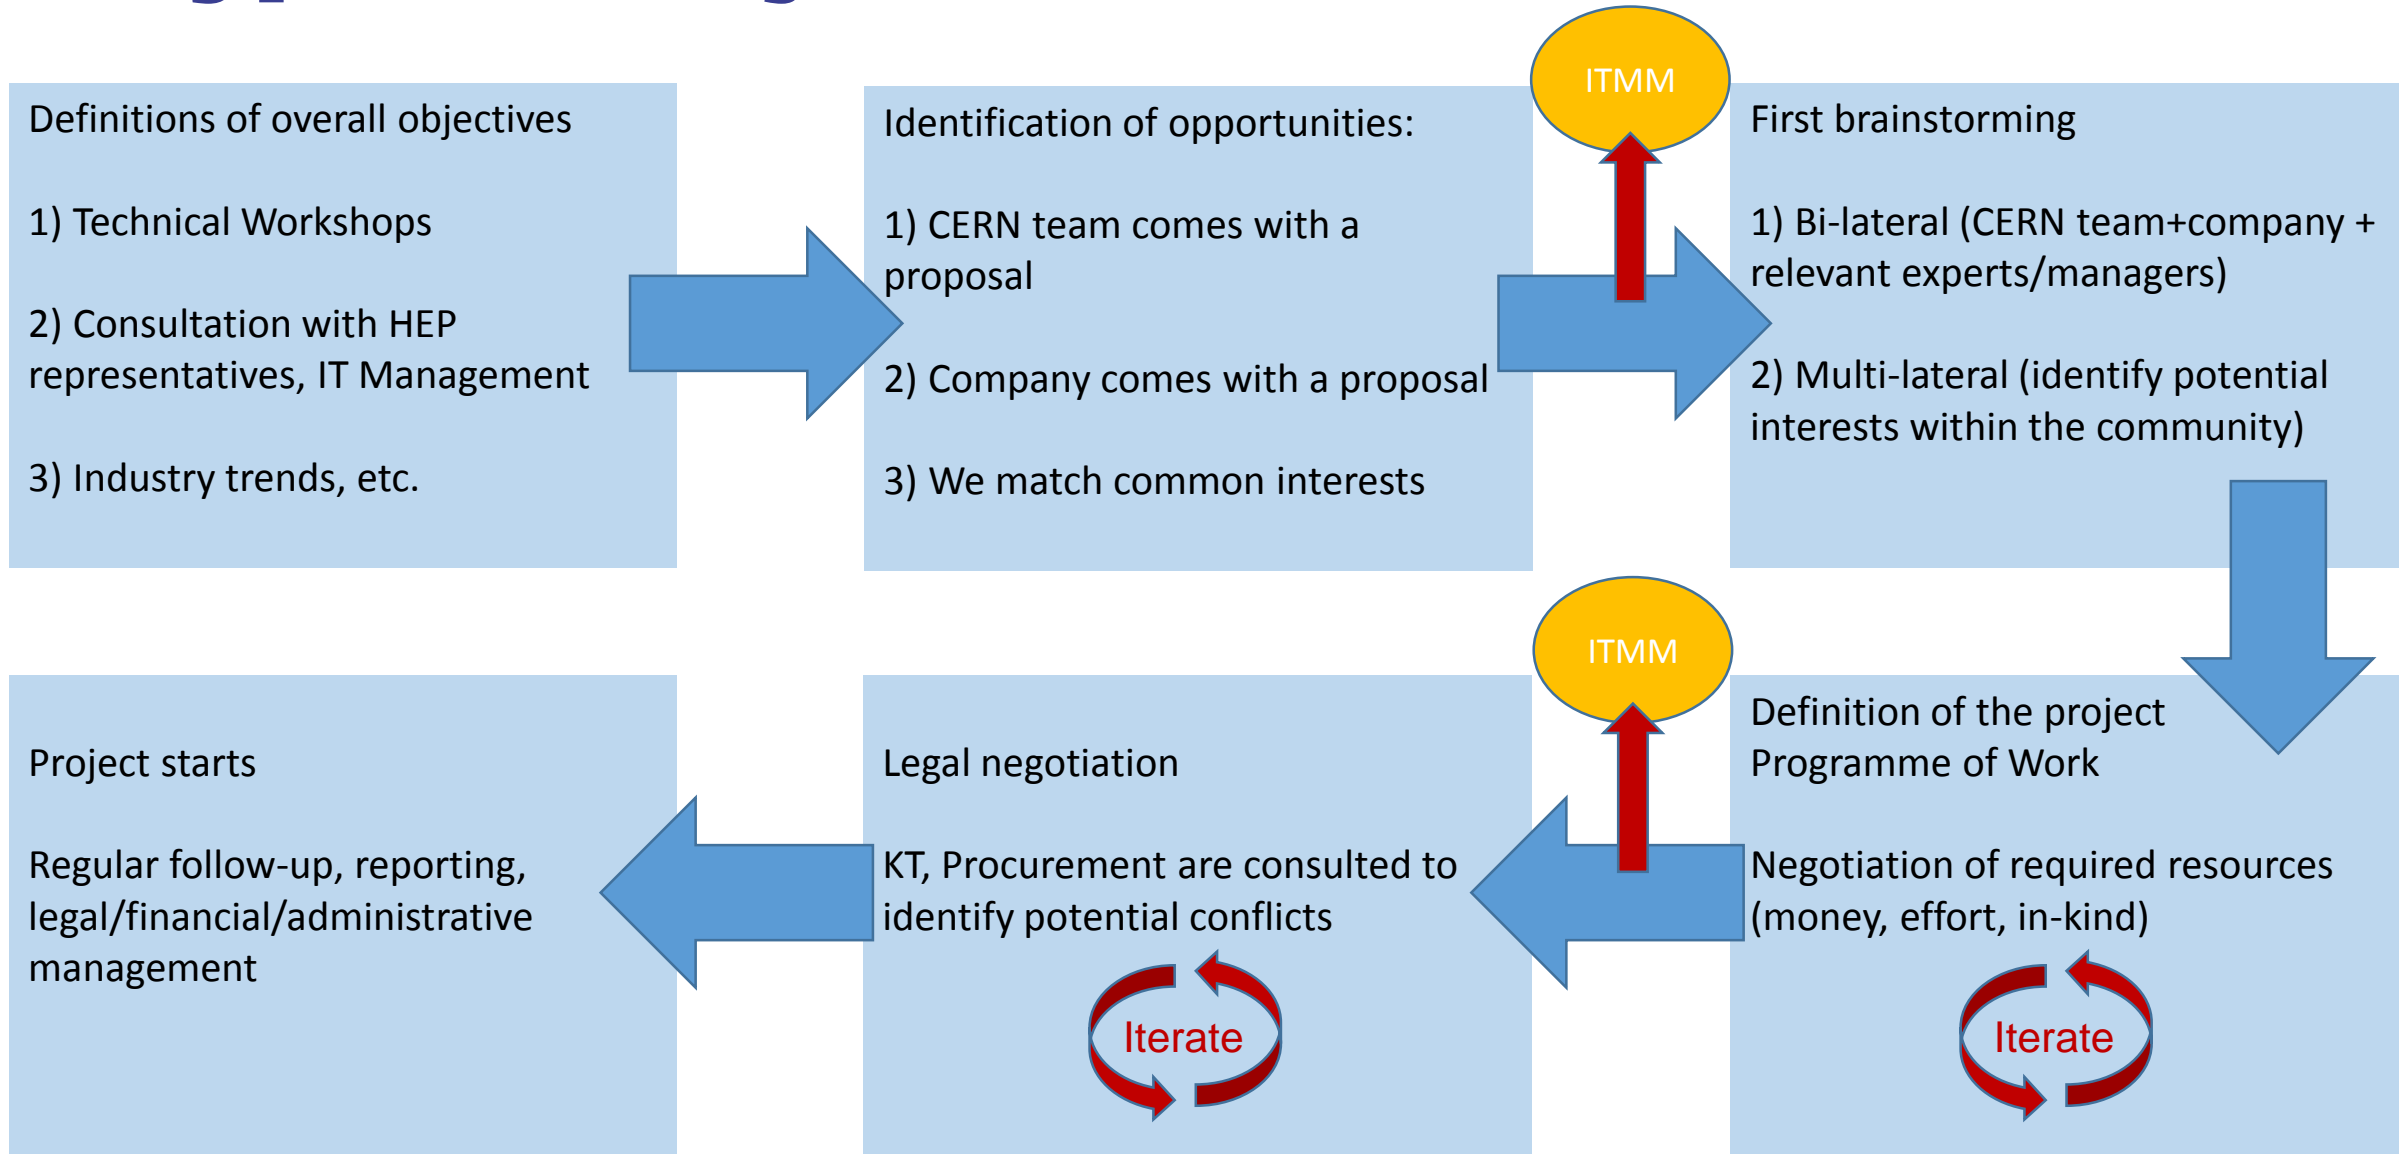
Communicate results, demonstrate impact, and reach new audiences.



Collaborate and exchange ideas with other communities to create knowledge and innovation.

Train the next generation of engineers/researchers, **promote** education and cultural exchanges.

Typical Project Definition Flow



CERN openlab Contributions in 2017

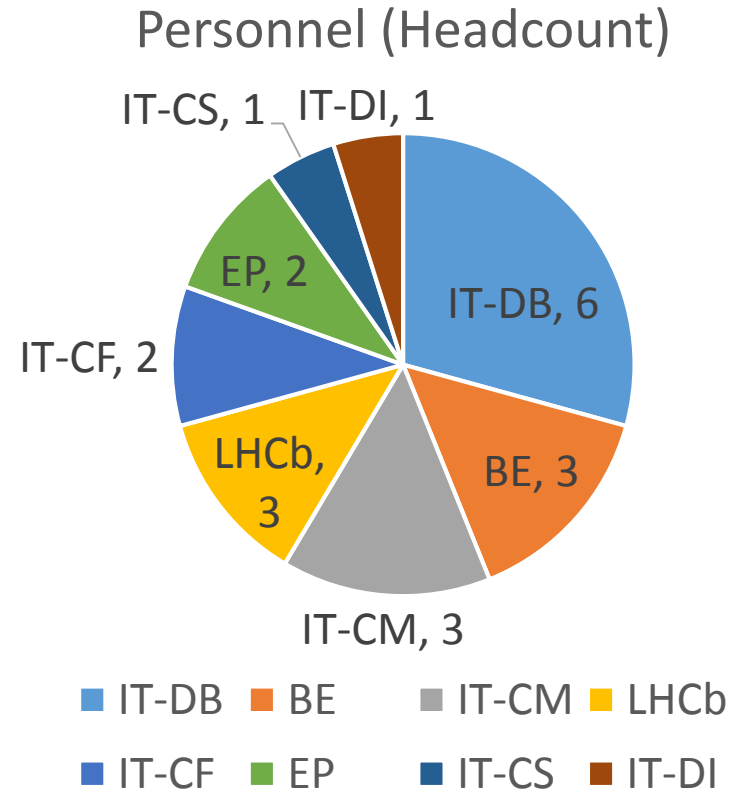
Funds to recruit ~20 FTEs (mostly Fellows)

Hardware and software (~250k CHF)

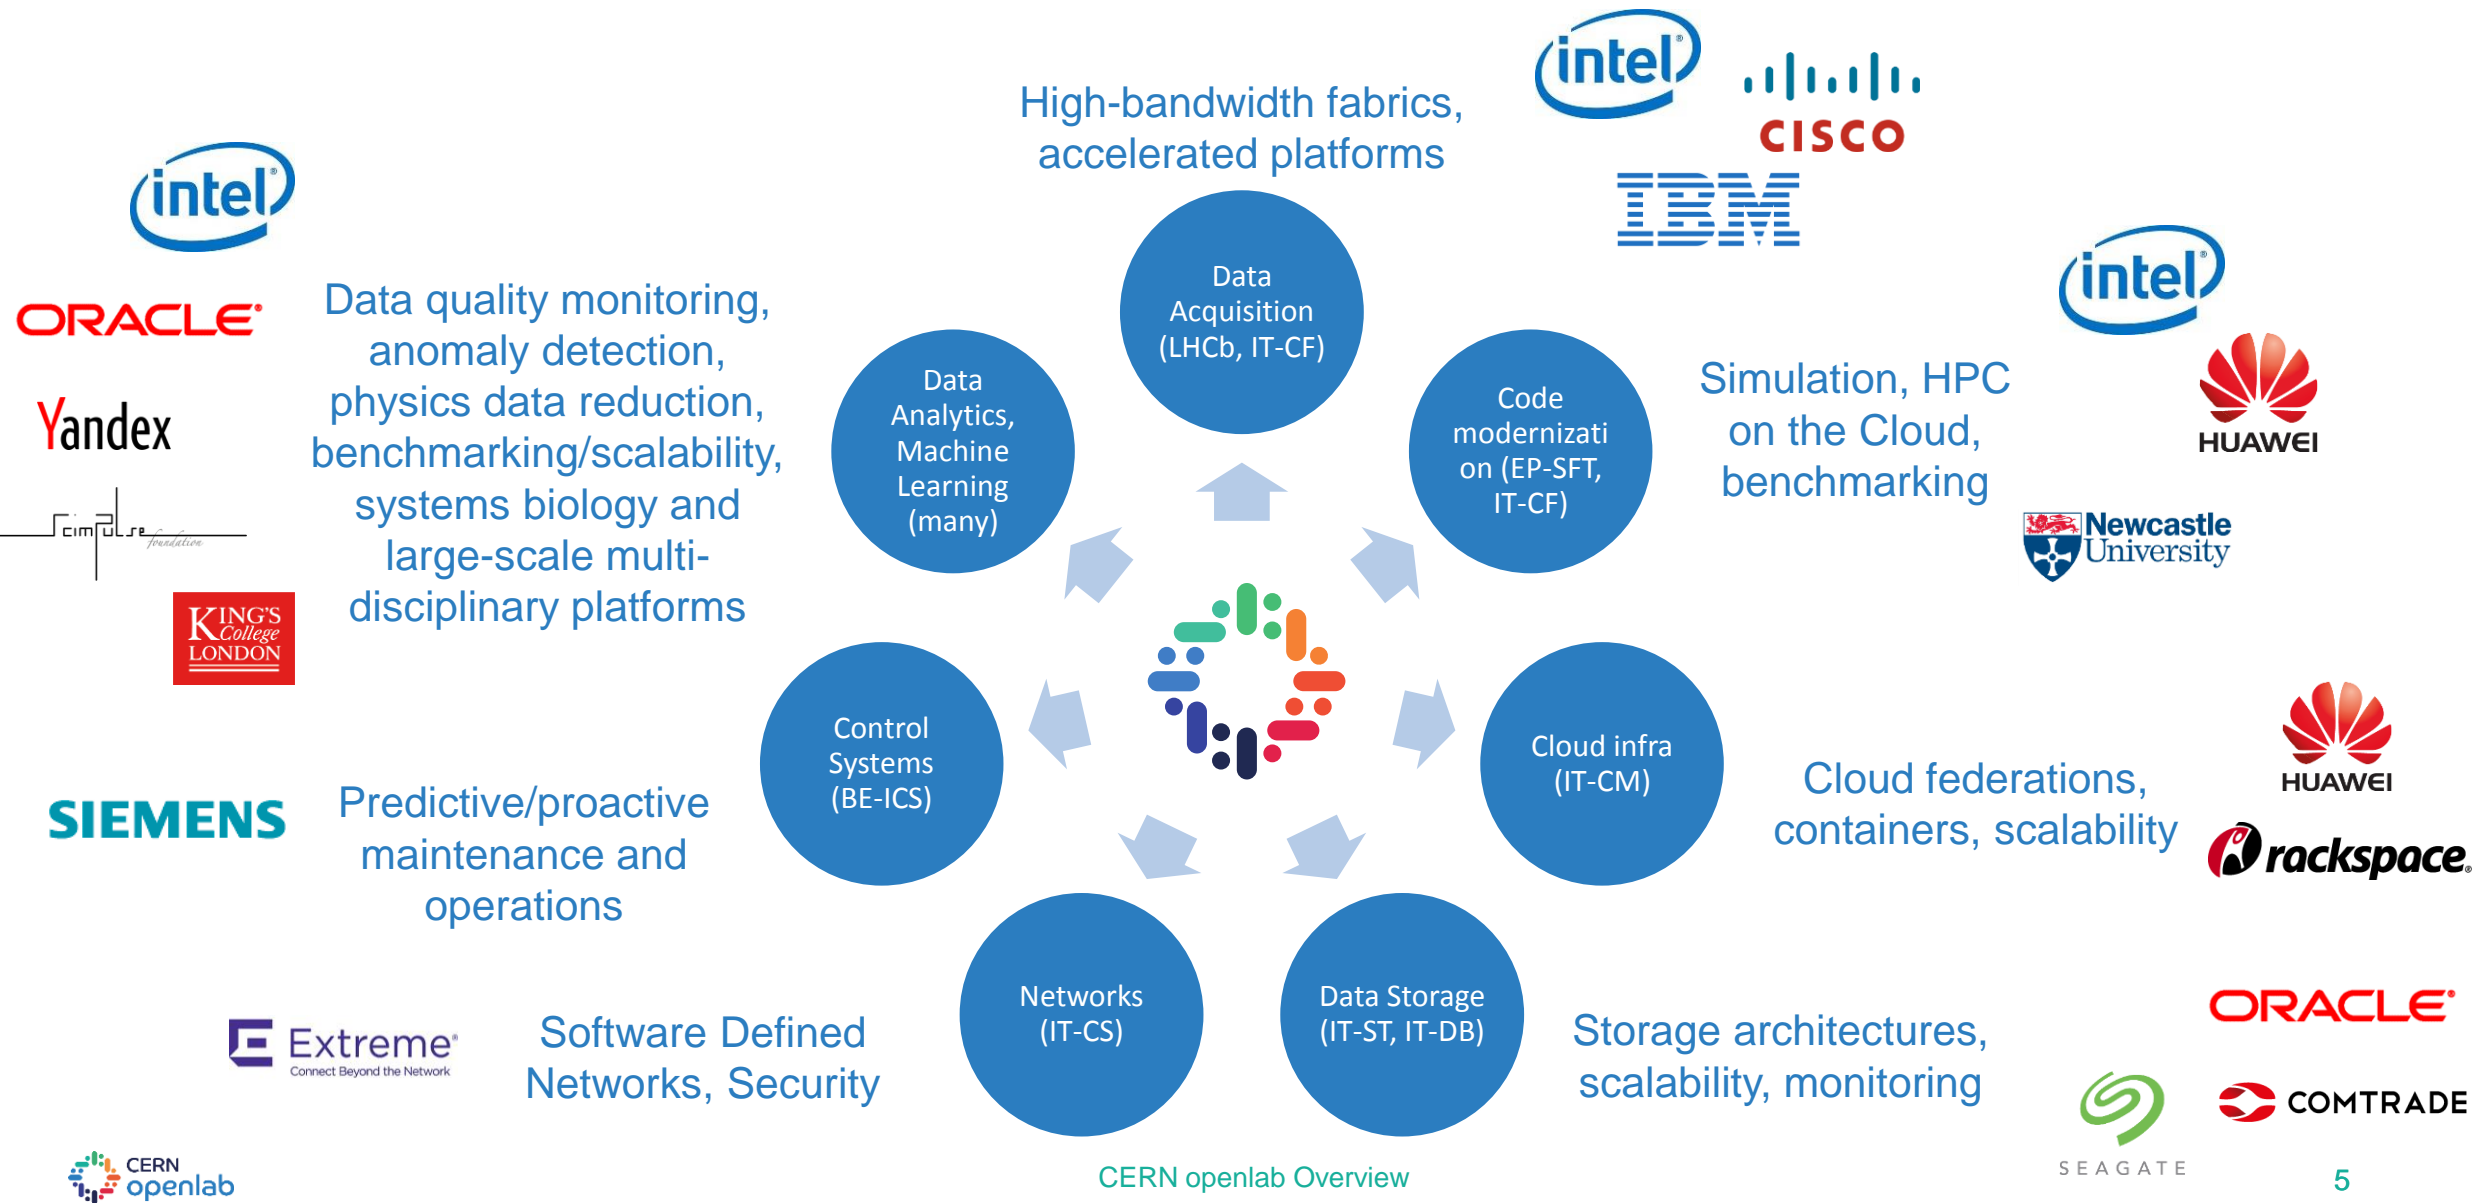
Expert technical support

Training and consultancy

Roadmaps, information events, etc.



JOINT R&D PROJECTS



Ongoing Projects

Who	Coordinators	What	How	When
Intel, IT-DB, Fermilab, CMS	Luca Canali (IT/DB) Oliver Gutsche	Investigation of Hadoop/Spark to accelerate data reduction in CMS	1/2 Fellow (Intel) The IT/DB Hadoop cluster Software licenses (Intel)	2017-2018
Intel, IT-DB, BE-ICS	Luca Canali, Fernando Varela	Machine Learning for LHC Controls log analytics	1/2 Fellow The IT/DB Hadoop cluster Software licenses (Intel)	2017-2018
Intel, CERN openlab	Federico Carminati	Machine Learning for Fast Simulation	IPCC grant Hardware and software (Intel)	2017-2018

Ongoing Projects

Who	Coordinators	What	How	When
Intel, EP-DT	Giovanna Lehman	Development of a fast, scalable K-V storage system for DAQ buffers	1 Fellow (Intel, EP/DT) Hardware (up to 2 full nodes in two years), software licenses (Intel)	2018-2019
Huawei, IT-CM	Tim Bell	OpenStack development	2 Fellows	2017-2018
Rackscale, IT-CM	TIM Bell	OpenStack development	1 Fellow	2017-2018
Extreme Networks, IT-CS	Edoardo Martelli, Stefan Stancu	Intelligent Bandwidth Optimization	1 Fellow	2017-2018
Comtrade, IT-ST	Luca Mascetti	EOS productization	Expertise	2016-2018

Ongoing Projects

Who	Coordinators	What	How	When
Oracle, IT-DB	5 different projects and coordinators	Data analytics, database technology testing, cloud platform services	5 FTEs	2018-2020
Siemens, BE-ICS	Fernando Varela, Filippo Tilaro	Data Analytics for LHC Controls	3 FTEs	2018

Planned Projects

Who	Coordinators	What	How	Duration/Status
Intel, IT-ST, IT-DB, Alice	Alberto Pace, Eric Grancher, Predrag Buncic	Assessment of 3D-Xpoint technology for in-memory applications	1 Technical Student/PJAS Hardware (2 full nodes)	2 years Technical write-up
Company1, Alice	Predrag Buncic, Latchezar Betev	K-V storage scalability and performance improvements	2 FTEs Hardware (up to 10 full nodes), software	2 years Technical write-up
IBM, CMS, LHCb, CERN openlab Possible participation of INAF	Maurizio Pierini Niko Neufeld Federico Carminati	Various applications of machine learning and software porting to Power 8+/Power 9 for DAQ, Data Monitoring, image analysis	1 DOCT 1 TECH (TBC) IBM Power 8+ cluster (plans to provide Power 9)	1-3 year Technical write-up

Planned Projects

Who	Coordinators	What	How	Status
E4, CMS, LHCb	Maurizio Pierini Felice Pantaleo Vincenzo Innocente Niko Neufeld	LHCb triggers CMS pixel tracking Fast inference, scalable training GAN-based fast simulation and analysis (detectors, image analysis)	1 PJAS at CERN 1 E4 engineer at CERN 4 engineers on-call Full nodes with different types of GPUs from low- to high-end Software	Up to 3 years Technical write-up
Intel, CERN openlab, INAF	Federico Carminati	GAN-based image analysis	1 FTE Hardware (Remote access and on- premise Nervana systems) and software	1 year Technical write-up

Planned Projects

Who	Coordinators	What	How	Status
Intel, University of Wisconsin	Miron Livny	Port and test HTCondor on newer Intel CPUs, test special sensors and instrumentation features	Hardware (Full Xeon-based nodes) and software	1 year Technical Write-up
Yandex, CMS	Virginia Azzolini	Data Quality and Popularity	Expertise, software	1-3 years Legal validation
Intel, Princeton, EP-SFT	Peter Elmer	Benchmarking of Root Vs. Other DA/ML tools	1 Fellow	1 year Legal validation
University of Eindhoven, SHiP	Eric Van Herwijnen	Development of Conditions Database	1 DOCT	3 years Signature

Planned Projects

Who	Coordinators	What	How	Status
Company2, CMS, Dune	Franz Meyers Emilio Meschi Maurizio Pierini Marzio Nessi	Investigation of FPGAs technology and frameworks for DAQ (real-time streaming Inference Engine for use in the Level-1 trigger systems)	2 FTE Hardware (FPGAs and boards) and software	1 year Technical assessment
Microsoft, ETHZ, CERN openlab, CMS, LHCb	Federico Carminati Maurizio Pierini Niko Neufeld	Investigation of FPGAs technology on-premise (DAQ) and in the cloud (training, simulation)	x DOCT Hardware (FPGAs and boards), cloud access	3 years Technical assessment
Company3, IT-DB, IT-CF	Alberto Pace, Olof Barring	Composable, low-cost SSD	TBD	TBD First contact


Suspended Projects

Who	Coordinators	What	How	Status
Huawei/Hisilicon, ATLAS, IT-CF, EP-SFT	Graeme Stewart David Abdurachmanov	Software porting and benchmarking on ARM64	28 nodes with various generations of Hisilicon ARM64 Technical support	Suspended, waiting for assessment of interest in ARM64 and manpower

Other activities

- **BioDynaMo**: application of simulation and cloud expertise to large-scale biological development simulations
- **GeneROOT**: application of Root to accelerate genomic analysis (in collaboration with EP-SFT)
- **LivingLab**: applications of machine learning and Natural Language Processing to medical data analysis and diagnostic support system (in collaboration with HSE)
- IoT applications for mobility, environmental monitoring (with SMB)
- Designing a high-level investigation programme around Quantum Computing
 - Programming models
 - Relevant applications in HEP
 - Expertise in cryogenics technology

SUMMER STUDENT PROGRAMME



In 2018

- 1840 applicants
- 40 selected students
- 14 lectures
- Visits to external labs and companies
- Lightning talks session
- 40 Technical reports

Summer Students Projects

Intel/EP-CMG-PS	Fast Inference on FPGAs for HEP trigger systems	Jennifer Ngadiuba, Maurizio Pierini
Intel/EP-LCB	Investigation of data direct I/O for 100 Gbit Ethernet	Niko Neufeld, Tommaso Colombo
Intel/EP-UAC	Deep generative models for calorimeter simulation	Stefan Gadatsch, Michael Kagan
NVIDIA/E4/EP-CMG-PS	HGCAL Fast simulation With Deep Learning	Shah Rukh Qasim, Jan Kieseler
IBM/EP-CMG-CO	Anomaly detection with machine learning for monitoring the quality of the data of the CMS experiment	Adrian Pol, Gianluca Cerminara
IT-DI-LCG	Quantitative Workflow characterization and modeling	Andrea Sciaba
IT-DI-LCG	Data Analytics and Machine Learning on Trident node monitoring tool	David Smith, Servesh Muralidharan
EP-SFT	Efficient unpacking of required software from CERNVM-FS	Jakob BLOMER, Gerardo GANIS
E4/EP-CMG-PS	Developing solution for large-scale network training and optimisation	Maurizio Pierini, Jean-Roch Vlimant
IT-CF/EP-CMG-CO	Benchmarking Machine Learning in HEP	Luca Atzori, Felice Pantaleo