

IT requirements for the next generation research infrastructures

CRISP Workshop: CERN - 1 February, 2013

Introduction

This document provides a summary of the “IT requirements for the next generation research infrastructures” workshop, a one day event held at CERN on February 1st 2013. This document does not provide full minutes of the meeting, only a brief overview of the sessions and their outcome, as well as the results of the post-event survey that was sent to all participants. The material presented during the event can be found on the following Indico webpage:

<http://indico.cern.ch/conferenceOtherViews.py?view=standard&confid=212402>

Objective

The objective of the event was to present the IT needs and procurement plans of the ESFRI research infrastructures from the larger physics domain that are working together in the CRISP FP7 project (<http://www.crisp-fp7.eu/>). The event aimed to bring together industrial representatives from the IT sector together with scientists/engineers engaged in IT and data management activities for physics, astronomy and analytical research facilities in Europe. A short summary of each session follows below.

Morning Sessions

Session 1: Presentation of the next generation of physics, astronomy and analytical research instruments and their common IT needs

Bob Jones opened the workshop by welcoming the participants and stating the objectives of the day. The first speaker was S. Bertolucci, CERN’s Director of research, who presented CERN’s achievements in 2012 but also talked about the immediate future plans for the increase of the LHC collisions’ energy to 14 TeV and further plans to increase the accelerator’s luminosity in 2018 and beyond. He discussed the exciting new Physics that the upgraded machine could unveil and also touched upon the IT and infrastructure requirements necessary to support the evolving scientific activities.

The second presentation by M. Krisch, the CRISP Project Coordinator, outlined the project's mission and scope and presented the CRISP Research Infrastructures. He went on to present the RI's upgrade plans and timelines and discussed the ways the RIs and industry currently work together and how the collaboration can be further explored.

During the third presentation L. Field, the IT & Data Management Activities Topic Leader, presented the challenges faced by the RIs. Using the example of ESRF he explained how the data is generated and discussed the variations in the requirements of the CRISP members. He then highlighted the synergies between labs, which are reflected in the four working packages: User Identity, Metadata catalogues, High-speed data recording and distributed data infrastructure. He pursued and described some of the work, challenges and collaboration in each of these areas.

The session ended with a networking coffee break.

Post-event participant survey evaluation of the first session: In general, the audience found the session interesting and gave a good high-level introduction to what the RIs are doing and what their challenges are, even though the plans could have been more detailed. The audience seems to have found it interesting to see the roadmap of big scientific projects and their requirements for computing. Overall, the audience found it gave an excellent overview, outlook and time-lines.

Session 2: The vision of IT industry

After the break, industry representatives presented their vision of the future and how they address the challenges put forward by the research community.

H. Cornelius, Technical Director for Advanced Computing presented Intel's vision of the Datacenter Imperatives for the future. He presented the evolution of datacenters with many divergent systems to a model where the datacenter as a whole is viewed as an integrated system providing IT services. He also discussed technology trends in chip design with Intel's plan to ship products at the 7nm scale by 2017, leading to more and more integration at the chip level. He touched on one implication of this in networking (Software Defined Networks) but also discussed the impact in terms of Big Data handling and HPC computing.

The second presentation was given by C. Clark, Director of Research in HP, on Software Defined Networks and how they will be faster, easier to manage and more optimized than today's networks. SDNs separate the control and data planes and logically centralize the control planes for the whole network allowing for optimization offered by a net-wide view. A third characteristic of such networks is the abstraction of the network infrastructure from the applications making them tolerant to changes in network technology. He also explained the openness and the flexibility of HP's approach to SDN.

The third presentation by J. Hughes, Chief Architect on storage for Huawei stirred-up much interest with a lively exposition of his vision and predictions for the future of storage technology. James explored some key concepts in the storage market where demand increases as prices go down. He went on to discuss the merits of technology concepts such as shingled disks which multiply the capacity but also the trend towards key-value stores in RAM, flash and disk. He also discussed object storage as a low cost option for online storage as a competitor to tape. Additionally, he made a number of interesting predictions about storage in the near future such as the replacement of

distributed file systems by scaled object stores, the complete replacement of fast disk storage by flash, the use of the key/value paradigm implemented in hardware for accessing RAM and the increased density of storage devices.

The last presentation of the morning by Jurry de la Mar from T-Systems gave a different perspective, discussing the challenges and solutions offered by a large service provider. Jurry stressed that it is essential to be technology neutral to satisfy a wide spectrum of customers but also to be able to respond quickly to customers' requests to scale operations within short delays. The structuring of the service portfolio in building blocks that can be combined together to a larger offering provides such flexibility for services. The use of the cloud as a deployment platform provides the necessary infrastructure with flexibility where he explained the company's participation in the *Helix Nebula – The Science Cloud* initiative (www.helix-nebula.eu). Jurry gave an example of a successful implementation with a big multinational and discussed the lessons learned. He went on to discuss T-Systems' research activities in setting up an HPC service for selected customers and investigating possibilities to run more services in a cloud environment. He touched on the idea of a federated cloud and the notion of the 'blue box' that would allow cloud services from different providers to be used seamlessly. Finally he pointed out the key issue of governance for such an interconnection, his firm belief being that it should be based on open standards and that it should result from a public-private partnership as the investment and governance issues are too big for any industrial player to take on individually.

Post-event participant survey evaluation of the second session: This session received very positive feedback and in general the audience found all talks to be interesting and well presented. The evaluation pointed out that this type of session could be even more interesting if the presentations were given by a person representing an RI and proposing an approach to tackling a given IT problem, and one from the industry proposing an envisioned solution. As one person pointed out; it is not always clear how the offer fulfills the RI requirements.

Afternoon sessions

Session 3: Procurement plans, Knowledge and Technology Transfer

The first presentation by J.F Perrin, head of the computing services at ILL, summarized the procurement plans of the CRISP Research Infrastructures and discussed the procurement process and the common practices between them. Jean-Francois went on to describe some of the challenges facing the RIs, such as the 'data deluge' and its impact on the infrastructure planning in the next few years, and emphasized that the RIs look for scalable long-term rather than 'big bang' solutions. He went on to stress the importance of stability and predictability in the relations with industry.

The second presentation by H. Mourin managing director of the Teratec initiative, a technopole dedicated to high performance simulation and computing, brought a different perspective on how research and industry work together using HPC to address issues covering a wide spectrum from Health & Energy to the Media. Teratec is driven by major industrial players. It also includes a large number of smaller technology providers but also works together with a few major research centers.

Teratec launched a number of joint projects which are managed by the Industry but also have extensive educational and training activities.

B. Jones, head of the CERN openlab, presented its activities during the third talk. The CERN public-private partnership recently celebrated 10 years of successful collaboration with Industry. Leveraging the LHC requirements, openlab provides a platform for future IT technologies, allowing Industrial partners to test and improve their products but also to develop ideas to meet the research community requirements. Bob emphasized the educational role of openlab in training the next generation of scientists and engineers. He finally talked about the collaboration with other research labs in expanding the model with the industry on the theme of 'Big Data'.

The last presentation was by M. Marinucci, Oracle's Director for the Research Industry. Announcing the 30th anniversary of the Oracle-CERN collaboration, Monica gave an overview of the Industry's perspective of the key components and drivers that both the Research community and the Industry have to consider if they are to engage in a successfully long term partnership. She went on and defined the area of common interest to both CERN openlab and Oracle, and the company's perspective on the partnership and how it translates to concrete benefits for its technology.

Post-event participant survey evaluation of the third session: The feedback indicated that this was a good session, worth attending, but slightly less inspiring. The openlab presentation was appreciated for giving a good insight on how a private-public partnership can work. It was stated that these topics deserve a break-out session, especially regarding procurement.

Session 4: Preparing a common roadmap

This session was organised as a panel discussion to highlight areas of common interest to the Research Infrastructures and the IT industry, gaps in market offerings, opportunities for joint development and next steps. The issues highlighted by the research Infrastructures (RIs) during the day were considered relevant to many business sectors. In particular, data interoperability for multidisciplinary research and automation of data centres are key issues for IT companies, especially when linked to the needs of data analytics in the realm of big data applications. Linked to the growing quantity of data that RIs will produce and manage, the issue of a long-term and low cost archival storage solution was viewed as a gap in the current market offerings. More data means more processing power required to analyse it and so the energy efficiency of the processors, storage and networking equipment will become a major factor in the cost effectiveness of future systems. Beyond technology, creating a pool of well-trained personnel who understand that data science is an area in which RIs and industry can work together. In terms of what are considered the best models for ensuring RIs and the IT industry can collaborate, a number of avenues were proposed. The CERN openlab model of a public-private partnership was considered suitable by the IT companies and RIs present. RIs and industry should consider a co-design approach whereby users become an active part of the creative development process of a product by interacting directly with the development teams. If the IT needs of the RIs can be sufficiently aligned, the approaches of pre-commercial procurement (PCP) and public procurement of innovative solutions (PPI: a step following pre-commercial procurement) could offer a means for creating a market opportunity with sufficient interest to encourage the IT industry to develop suitable products. The European Commission intends to encourage the uptake of PCP and PPI in Horizon 2020 as a means of reducing the gap between research and market commercialization.

Post-event participant survey evaluation of the fourth session: The panel was well rated for being open and interesting, and for discussing relevant points. Some survey responses suggested it would be more useful to focus on concrete examples and to answer the "HOW" question. It was proposed that it would be more interesting if a preliminary agenda was prepared, with a limited number of questions to discuss and the commonly agreed answers can be selected at the meeting and fed into a roadmap or plan for the future.

Further results of the evaluation

The audience appreciated the one-day format and the majority would be interested in attending again (95 %). Many replies pointed out that the morning sessions can be broad, but it is preferable to have focused talks or even parallel sessions in the afternoon. One person who answered would like to see a two day workshop with RIs presenting their IT needs, problems and questions while on the second day industrial people could present possible solutions. The RI presentations should be sent to the industrial presenters before the event to ensure they can consider possible solutions. Overall the feedback indicated that a future workshop should be more focused on specific topics.

Potential IT Technologies subjects

Processing	16%
Storage	26%
Networking	16%
HPC	12%
Cloud computing	30%

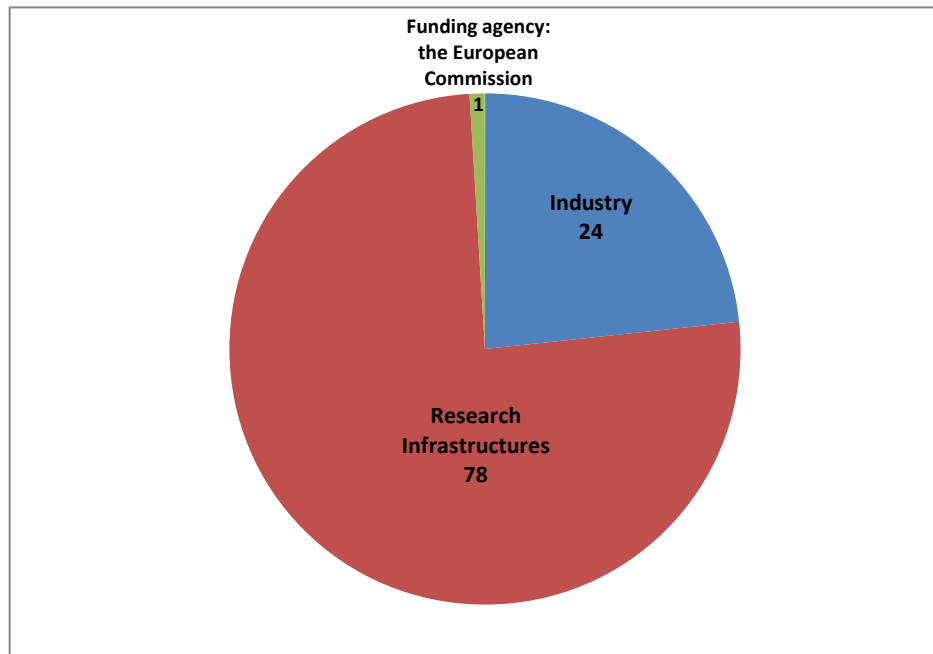
Other topics

Funding and business models for co-development of IT solutions	21%
Mobility of personnel and training	10%
Identity management	23%
Big data (data analytics)	21%
Long term archiving	15%
Software licensing and open source developments	10%

To summarize, 62% of the participants found the workshop interesting, 38 % very interesting. There is a pronounced interest to attend an event of this kind in the future.

Distribution of participants (103 Total)


IT requirements for the next generation research infrastructures



This high-level event organised by the CRISP consortium was streamed live on 01 February 2013 from 09:00 to 17:00 CET.

The webcast had 781 single IP connections – with a record 307 IPs connected at once. The total connections can be broken down by countries:

Visitors Geography



The high number of connections from Switzerland can be attributed to CERN personnel watching the event on-site.