

Offering Global Collaboration Services beyond CERN and HEP

J Fernandes¹, P Ferreira², T Baron³

¹ CERN, 1211 Geneva 23, Switzerland

² CERN, 1211 Geneva 23, Switzerland

³ CERN 1211, Geneva 23, Switzerland

E-mail: Joao.Fernandes@cern.ch

Abstract. The CERN IT department has built over the years a performant and integrated ecosystem of collaboration tools, from videoconference and webcast services to event management software. These services have been designed and evolved in very close collaboration with the various communities surrounding the laboratory and have been massively adopted by CERN users. To cope with this very heavy usage, global infrastructures have been deployed which take full advantage of CERN's international and global nature. If these services and tools are instrumental in enabling the worldwide collaboration which generates major HEP breakthroughs, they would certainly also benefit other sectors of science in which globalization has already taken place. Some of these services are driven by commercial software (Vidyo or Wowza for example), some others have been developed internally and have already been made available to the world as Open Source Software in line with CERN's spirit and mission. Indico for example is now installed in 100+ institutes worldwide. But providing the software is often not enough and institutes, collaborations and project teams do not always possess the expertise, or human or material resources that are needed to set up and maintain such services. Regional and national institutions have to answer needs, which are growingly global and often contradict their operational capabilities or organizational mandate and so are looking at existing worldwide service offers such as CERN's. We believe that the accumulated experience obtained through the operation of a large scale worldwide collaboration service combined with CERN's global network and its recently-deployed Agile Infrastructure would allow the Organization to set up and operate collaborative services, such as Indico and Vidyo, at a much larger scale and on behalf of worldwide research and education institutions and thus answer these pressing demands while optimizing resources at a global level. Such services would be built over a robust and massively scalable Indico server to which the concept of communities would be added, and which would then serve as a hub for accessing other collaboration services such as Vidyo, on the same simple and successful model currently in place for CERN users. This talk will describe this vision, its benefits and the steps that have already been taken to make it come to life.

1. Introduction

The LHC and HEP community at large has had always very a special and unique set of requirements regarding the collaborative tools environments. These requirements have been present since the middle 90s, much before other research communities or commercial organizations, due to the nature of the HEP collaborations: highly dispersed across the globe producing both an important volume of



document data and real-time communication traffic. As a result, CERN had to build an integrated framework of collaborative tools, to be able to cope with these.

The ecosystem built has been considered a success by several organizations, specifically the ones that over time, started to have similar requirements as the LHC ones. Following this success, several research and education organizations have requested the same level of services to be provided by CERN to their own communities.

This paper intends to describe the efforts that are taking place to enable these service extensions and the vision of how the models proposed will evolve, fully aligned with CERN's mission.

2. Current Services Status

2.1. Vidyo

Since 2012, Vidyo at CERN relies on a globally deployed infrastructure, as depicted below in Fig. 1.

The service is now consolidated with an intensive usage on the day-to-day activity of the community.




Fig 1. – CERN Worldwide deployed infrastructure

The large-sized user base has multiplied the number of use cases. The service is used from regular daily working meetings, to remote teaching or outreach events. The current usage is now reaching 8000 meetings per month over more than 20000 users registered, translating in more that 50M minutes of videoconferencing per year.

2.2. Indico

Indico is an open source event management system, developed at CERN. It was first conceived as a conference management system but later evolved into a full-fledged event organization solution with a special focus on collaborative tools at the organizational level. It seamlessly integrates with Vidyo, allowing video-conferencing meetings to be easily scheduled and configured.

In recent years, Indico usage has consolidated itself, as showed in Fig 2., within the High Energy Physics community – currently, more than 150 different Indico installations exist around the world.




Fig. 2. – Number of Events managed by Indico CERN.

If this increase in usage has been of great benefit for many research communities around the world allowing them to conduct their own events in an independent fashion, the truth is that demand for CERN's Indico server has continued to grow. Empirical evidence shows that HEP users tend to prefer Indico at CERN because of the central role that it plays, as the organization itself, in this particular research field. Most HEP users are in some way connected to CERN and tend to have a CERN account. At the same time, CERN's Indico is the first server in the world to receive every new version and enjoys a "Premium" level of support due to the simple fact that Indico is developed in-house and the development team provides technical support in case of technical issues and bugs.

3. Motivation for Expansion

The fact that both Vidyo and Indico services and the integrated workflow that they provide enjoy a high degree of satisfaction within their HEP user base, not only in terms of user experience but in what concerns reliability and quality of support, has triggered the interest of several organizations within and beyond HEP. The performance of both systems in several aspects that include service integration, scalability, stability, criticality, among others serves as an example to follow to many communities around the globe, which share a common set of requirements but, at the same time, are limited in terms of operation, due to their national or regional mandates (that often contradict their *de facto* capabilities). In addition, very often, institutes, collaborations and project teams do not possess the expertise or human/material resources needed to set up and maintain such globally available services. All these reasons fomented the approach to CERN with a basic fundamental question: how can such organizations use globally deployed services such as Vidyo or Indico, in a secured, dedicated and customized way, seamlessly and fully integrated with the communities that they are supposed to serve?

The answer to this question needs the establishment of a new strategy, model and underlying architecture. One could propose to "use the services as they are", but the issues in such approach are several: for example, the user base would always need a CERN account to create or moderate meetings and remote sessions; some of the users do not possess a CERN account, as their research is not connected directly to HEP. Another example is meeting management across several instances of the same tool: in Indico for example, the duplication of meetings with the same topic on different servers is not at all desired. It would highly decrease efficiency, as there would be no "unique place" where to find meeting materials or access events.

At the organizational level there are also costs: a local infrastructure needs to be maintained, implying support fees and manpower. In the event of a global centralized offer, most of those "collective resources" could be optimized and part of them potentially liberated and assigned to other areas where they could be put to a better use.

In addition, CERN's services and expertise is appealing, representing important advantages: a centrally supported infrastructure, already in place, running the most up-to-date versions, managed by experts in both systems with support mechanisms already in place with consolidated processes of continuous service improvement.

4. Model Proposed

In order to make a global offering, eliminating the issues mentioned, a new model has been proposed. The proposal consists on a worldwide service hosted and maintained by CERN, offered as a single but distributed “Infrastructure as a Service” (IaaS) model.

The CERN IaaS can be then shared between the supported communities, but properly isolated with possibilities of customization on a per community basis. The access is controlled and relying on federated user identity models, in agreement and fully compatible with each community AAI (Access, Authorization and Identification) schemas and policies in place.

The model has been developed on three main components: technical, cost recovery and legal. The technical model, being the first component to be conceptualized and drafted is showed in the picture below.

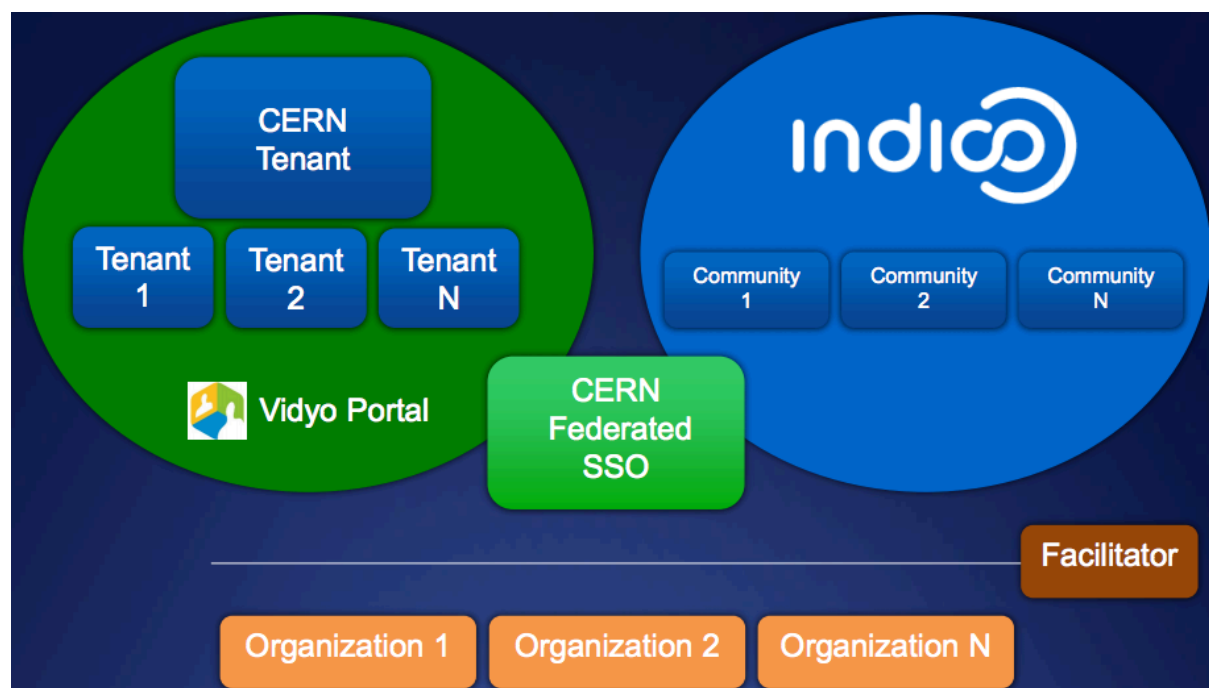


Fig 3. – Technical Model proposed for service extension.

In order to materialize all the three components, CERN has started a partnership with GÉANT Association.

GÉANT is the Pan-European research and education network that interconnects Europe’s National Research and Education Networks (NRENs). The GÉANT network connects over 50 million users at 10,000 institutions across Europe, supporting research in HEP but also in areas such as energy, the environment, space and medicine.

Even if the main affiliates are based in Europe, it possesses a set of privileged relationships and collaborations with several stakeholders, both scientific and commercial across the globe. The nature of the organization makes it the ideal partner in order to put in place such an effort of service globalization for the worldwide Research and Education (R&E) field at large.

In the framework of this initiative, GÉANT will act as a facilitator basically for two main reasons: in one hand, it’s an organization very close to potential “end customers” as a network provider for science, research and education; on another hand, due to the nature of its activity, it has already in place all the necessary mechanisms to charge costs and user support structures. These aspects are important for the cost recovery component of the model. The base idea is not to make profit from such an offer (also due to the nature of organizations such as CERN and GÉANT) but the recurring costs of such an expansion need to be recovered. At the same time, the cost recovery model needs to be

balanced, in order to avoid any potential disruption of the ongoing commercial market, providing grounds for unfair competition, arising potential legal liabilities.

Finally, the legal component of this model has also been proposed. As a first initiative of this kind within CERN, it has been the aspect of the described initiative that required more deep work to setup. Several aspects as the already mentioned potential unfair competition, commercial market disruption and compatibility with the CERN's governing rules and regulations, and general mission of the organization, required the advice and expertise of the CERN Legal service.

5. Future Challenges

Vidyo and Indico are now consolidated and heavily used services within the HEP community with a high rate of user satisfaction, success providing an overall excellent end user experience. This success has triggered a lot of interest from communities within and beyond HEP to profit from these services.

As a result of this demand, CERN has put in place a model to be able to extend in a controlled, secured and scalable manner, the offering of these services to the worldwide research and education communities.

In 2015, one should start to witness the first results of this initiative. A first pilot for a European NREN is starting by middle in June. Taking into account the interest shown, other organizations may rapidly follow. In a more technical aspect, and specifically in Indico, new features as the concept of community and search within a given community will need to be made available.

If one anticipates 2016, a phase of service consolidation will immediately follow. Aspects like improvement of AAA models (basically AAI extended with Auditing and Accounting), also because of several initiatives as EduGain, that are rapidly evolving and can be a very important component within this effort. In another hand, as the service evolves and gets wider interest, specially for cases of organizations where Vidyo and Indico will be provided integrated, the development of cross service business models will also need to be done.

Acknowledgements

The authors would like to acknowledge and thank the CERN IT Department and the Experiment Management and Representatives for supporting all the work done in the field of Collaborative Tools for the LHC. In addition, the author would like to thank the IT-CIS Group Leader and CERN IT DHO, for all the support and guidance.

References

- [1] J Fernandes, T Baron, B Bompastor, Experience from the 1st Year running a Massive High Quality Videoconferencing Service for the LHC, 20th International Conference on Computing in High Energy and Nuclear Physics (CHEP2013) IOP Publishing Journal of Physics: Conference Series 513 (2014) 062014 doi: 10.1088/1742-6596/513/6/062014
- [2] Gonzalez Lopez, J B and Ferreira, J P and Baron, T 2010 Indico Central - Events Organisation, Ergonomics and Collaboration Tools Integration, J. Phys.: Conf. Ser. 219 (2010)082002, doi:10.1088/1742-6596/219/8/082002
- [3] Indico, <http://indico.cern.ch>
- [4] CERN Vidyo Service Website <http://information-technology.web.cern.ch/services/fe/vidyo/>
- [5] Vidyo Service Dashboard Website <http://avc-dashboard.web.cern.ch/Vidyo>
- [6] IaaS, PaaS, SaaS (Explained and Compared) <https://apprenda.com/library/paas/iaas-paas-saas-explained-compared/>