



HORIZON 2020 FRAMEWORK PROGRAMME

IOStack

(H2020-ICT-2014-7-1)

Software-Defined Storage for Big Data on top of the OpenStack platform

Communication strategy and plan

Due date of deliverable: 31-12-2015
Actual submission date: 31-12-2015

Start date of project: 01-01-2015

Duration: 36 months

Summary of the document

Document Type	Deliverable
Dissemination level	Public
State	v1.0
Number of pages	10
WP/Task related to this document	WP6/T1.1
WP/Task responsible	MPS
Author(s)	Raúl Gracia-Tinedo, William Oppermann, Michael Breen
Partner(s) Contributing	URV, IBM, BSC, EUR, IDI, ARC, GRP
Document ID	IOSTACK_D6_1_Public.pdf
Abstract	In this document we define the required process and strategy for dissemination activities. We also provide a description of the dissemination and community involvement activities achieved until M12.
Keywords	Dissemination, Communication, Community

Table of Contents

1	Executive Summary	1
2	Process and strategy for dissemination activities	2
3	First Year Dissemination Activities	4
3.1	Academic conferences and publications	4
3.2	Mass Media and Digital Media	5
3.3	Industrial Dissemination events	6
3.4	External collaborations	7
4	Community Involvement and Exploitation Activities	7
5	Future Dissemination Activities	8

1 Executive Summary

European research projects are meant to make a positive impact on the society. Inherently, this yields to let the general public know and understand the objectives and achievements of a project. For this reason, the IOStack project aims at disseminating its results to the community and spread the word about its novel advances related to Software-Defined Storage (SDS) and Big Data technologies.

To properly address this goal, we designed a communication plan for IOStack. The plan consists of a battery of guidelines and dissemination actions that each partner in the project should follow. As a result, IOStack will maximize the visibility and impact of its results.

In this sense, as the project involves both industrial and research partners, the impact of dissemination may consider actors in both industry and academia. This means that dissemination actions should focus on their targets to get their attention, so we could assure a better broadcasting of the results. We will extend this point throughout the document.

Apart from the communication plan, we will summarize all the dissemination activities done by each partner. Despite in the first year most of the work is focused on research and development, the consortium has already achieved several dissemination actions that spread the word about the IOStack project internationally. We divided this activities into: i) academic conferences and publications, ii) mass and digital media impacts and iii) industrial dissemination events.

Finally, we included in this document all the external collaborations and the community involvement activities to test some of the first year results achieved by partners.

2 Process and strategy for dissemination activities

In this section, we provide a general overview of the dissemination plan in IOStack including both *guidelines* and types of *dissemination actions* to be performed.

Our first guideline for dissemination is to *specialize dissemination events* to reach well-defined user and practitioner communities. Therefore, in order to communicate properly and not to overlap themes between partners, partners should focus on their own Work Packages (WPs), which are related to a specific topic in the project. By doing this, each partner can be specialized on particular aspects of Software-Defined Storage (SDS). Such a degree of expertise and specialization will provide high value to the dissemination actions and communication activities of partners. Furthermore, partners should benefit from their specialized background to infer the best alternatives to disseminate their own WP results, maximizing the visibility of IOStack and its impact to society.

We distinguish four specific topics that will drive the communication plan:

- SDS for object storage
- SDS for block storage
- Compute/virtualization orchestration
- Cooperative/active storage & data analytics strategies

As a second guideline, since the IOStack project relies on an association of industrial and research partners, it is very important to take advantage of this plurality for *broadcasting each novel development in both communities*. To do so, each partner should disseminate its results using suitable channels for each action. It means that research partners should *publish research papers* and participate in events or workshops and for the industrial partners they should send *press releases to mass media* and spread the word of early results in local or international events related to Software-Defined Storage (SDS) technologies.

In third place, apart from the single communication plan of each partner and its dissemination actions, *joint collaborations among partners and other research projects or industries* are also important, because they could result into a new outcome. Furthermore it is an opportunity to propagate the achievements to other stakeholders.

Apart from these high-level guidelines, in IOStack we define a battery of specific actions to materialize the project dissemination. To this end, Fig. 1 describes the different types of dissemination events considered in this project. In the following, we describe their objective and how they can promote the visibility and adoption of IOStack.

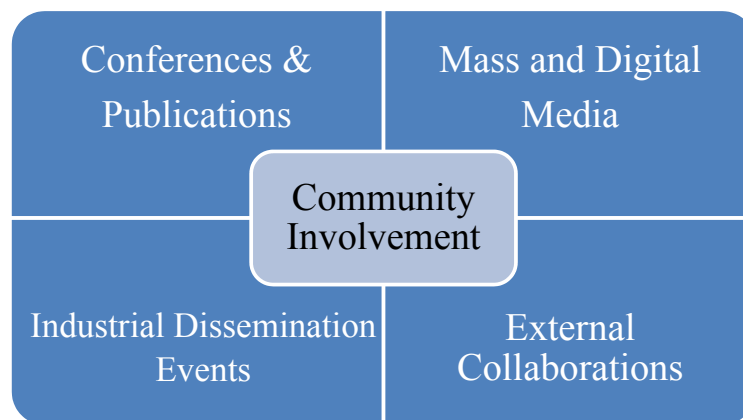


Figure 1: Areas of our dissemination plan for each topic/aspect of IOStack.

Conferences & Publications: Research results that have been achieved in the different work packages will be presented to the scientific community in the form of publications in journals, presentations and demonstrations in conferences, workshops and exhibitions. This is key to promote the ideas behind IOStack, and to have the scientific backing for introduction in the relevant standards later on. The research quality of the academic partners and the good synergy with industrial ones should ensure high impact publications in top conferences and journals.

Some of the conference where we are planning to attend are: IEEE Cloud, IEEE P2P, ACM Middleware, ACM CCS, ACM Eurosys, USENIX FAST, among others. Regarding journals, we expect to publish in some of the following ones: Elsevier Future Generation Computer Systems, Elsevier Computer Networks, Elsevier Computer Communications, IEEE Internet Computing, IEEE Network Magazine, IEEE/ACM Transactions on Networking, VLDBJ, and ACM/IEEE Transactions on Database Systems.

As described in deliverable D2.1, these published results will follow a *green open access* strategy where articles and related data will be self-archived by the authors in an online repository (project's webpage and/or personal homepage) after the corresponding scientific publisher grants open access to authors. Most conferences and journals allow publishing the results after the peer-review period as long as the final copy-editing of the article is kept to subscribers or attendees of the journal and conferences respectively.

Mass and Digital Media: The IOStack partners will engage in the usual forms of dissemination to large audiences. Today, this is clearly related to mass and digital media, which includes channels such as the Internet, TV and radio, among others. We should take into account that this form of dissemination is general and broad, as it is not specially tailored to a technical audience. This yields that we should carefully craft the contents of these mass dissemination events in order to be easily accessible for non-technical audiences.

Therefore, IOStack partners will develop the required collateral to promote the adoption of the IOStack project including social media videos, user manuals and marketing collateral. We plan to expose some of these materials to the commissioners in the next project reviews in order to acquire high-quality feedback.

Industrial Dissemination Events: This category of dissemination contain those events that are neither research nor mass/digital media events. For instance, industrial meetings or technical/developer conferences fall into this category. In fact, category of dissemination activities is of great importance in IOStack; it includes, for instance, *OpenStack Summit* or *DockerCon* events, which are premier technical meetings for the open-source community that supports the technologies behind IOStack.

During the project, assisting to such events (OpenStack Summit, DockerCon) will be a dissemination priority. This will make the community interested in IOStack, which may lead in contributions from the open-source community to the project. Ideally, we expect to up-stream the code of IOStack as an official part of the OpenStack ecosystem.

External Collaborations: In our experience with other EU projects, to weave alliances with other research project, industries and initiatives will be a key element to propagate the impacts of IOStack rapidly. Moreover, the influence of partners within a collaboration may enable a virtuous circle in which all collaborators are benefited. The mutual benefits of collaborations among projects includes the discovery of new use cases, the inter-operation of various EU software toolkits, and the increment of user/customer communities, to name a few.

For these reasons, we will also focus on other European Projects in the second year of IOStack. As we will mention later on, we started several contacts with NESUS initiative and various companies from the energy utility sector during the first year, so we are planning the activities to develop with all of them. We expect astonishing results to come in the next months.

Community Involvement: In IOStack, we target 3 different communities: *datacenters*, *data analytics companies* and *open-source communities*. To reach these communities will boost the adoption of IOStack

by real-world companies. Moreover, as well as it will help further developments of the project toolkit beyond the project life time.

First, we target the involvement of datacenters once the IOStack prototype is ready. Concretely, during the second year of the project, we will resort to the partners that operate datacenters (MPStor, Arctur) in order to make the IOStack prototype a usable tool for administrators. In this way, we aim to disseminate our achievements and to gather the feedback and interest from the participants.

Second, we also aim at involving our data analytics use cases to do pilot tests with IOStack. In particular, we are currently developing SDS services for improving, both in terms of flexibility and performance, the data analytics tasks of Idiada and GridPocket. Once achieved our objective, we need to publicize these success stories to other other companies, so that they may also adopt IOStack in their own business processes. Clearly, appropriate industrial dissemination events and venues are key to spread these success stories to increase our potential customer community.

In addition, open-source community involvement will be very important for the project too. There is another interesting community that we will involve in the IOStack project. We are pointing to the OpenStack community, because OpenStack is becoming the “de facto” world standard as a cloud storage platform, with partners such as: NASA, RackSpace, Intel, Paypal, CERN and many more. As IOStack is supported on this platform, we plan to involve the OpenStack community in order to get their feedback and seek valuable collaborations around the globe. Similarly, the Docker community will also be very important for the compute framework of IOStack. We will also execute specific actions to attract the Docker community to the technical development of the project.

Finally, for the third year we plan to consolidate the activities started in the second year, seeking a major impact in industrial field, enforcing the academic research results and finally getting a bigger user communities involvement. Depending on the outcomes achieved in the second year, we will adapt the dissemination activities in order to take the best advantage of these results. It means that getting a positive feedback of the user community could help us define the best way to market our achievements and collaborations with other EU projects could derive to new goals into the IOStack project. In summary, we can define some guidelines for the third year goals and activities, but we will be flexible on that point and set up a fine plan by the end of the second year.

3 First Year Dissemination Activities

In this section we will summarize the first year dissemination activities, specifying the partners participating in each activity. We divided them into the following subsections:

3.1 Academic conferences and publications

- **Conference paper** (February 16-19, 2015, Santa Clara, CA, USA): “SDGen: Mimicking Datasets for Content Generation in Storage Benchmarks”. Authors: Raúl Gracia-Tinedo, Danny Harnik, Dalit Naor, Dmitry Sotnikov, Sivan Toledo and Aviad Zuck. *13th USENIX Conference on File and Storage Technologies (FAST’15)*[1].
- **Magazine paper**: “IOStack: Software-Defined Object Storage”. Authors: Raúl Gracia-Tinedo, Pedro García-López, Marc Sánchez-Artigas, Josep Sampé, Yosef Moatti, Eran Rom, Dalit Naor, Ramon Nou, Toni Cortés, William Oppermann, Pietro Michiardi. *IEEE Internet Computing* (to appear)[2]. This paper will be published in a high visibility magazine. This publications has a twofold benefit: First, a magazine enables us to spread the ideas of IOStack to a broad audience; and second, the article offers a general and coherent view of the project, as all partners have participated in its elaboration.
- **USENIX FAST’15 Conference** (February 16-19, 2015, Santa Clara, CA, USA): URV and IBM presented a full research paper related to realistic data generation for storage benchmarking, specifically in terms of compression. Apart of being a joint work between two partners of the consortium, it was a good opportunity to promote and augment the visibility of the IOStack project though face-to-face meetings with several attendees.

- **2nd URV Doctoral Workshop in Computer Science and Mathematics (Tarragona, November 2015):** Pietro Michiardi from Eurecom was invited as keynote speaker at the 2nd URV Doctoral Workshop in Computer Science and Mathematics (November 13th 2015). Apart of being an interesting joint dissemination activity between Eurecom and URV, the talk of Pietro Michiardi exposed the scheduling and allocation problems currently found in many industrial and real world compute environments. This is directly related to the challenges to be solved in IOStack. More details can be found at the workshop's webpage¹.

3.2 Mass Media and Digital Media

To be present in digital media is key to disseminate the results of IOStack. In this sense, our actions will be focused both on publishing the news related to IOStack and to participate in mass media for a broader impact on less technical audience (tv, radio, etc.).

Regarding digital media actions, in IOStack we already executed the following actions:

- **IOStack official webpage²:** We set up an official project webpage. The IOStack site offers general information about the project objectives, the consortium and the software and dataset outcomes. Moreover, the IOStack webpage offers means to specialized and general public of auditing the progress of the project via public deliverables (open data). The webpage contains a news section to link the most important events related to the project, as well as integration with IOStack Twitter account.
- **IOStack Twitter account³:** IOStack is present in Twitter as a mean of rapidly disseminating news related to the project. The account is being updated frequently and it is being linked with other accounts of previous EU projects, in order to maximize the impact of IOStack tweets.
- **IOStack Github account⁴:** Another way to disseminate the project is via promoting our source code outcomes. For this reason, we are actively working in the IOStack Github account that is publicly available.
- **IOStack links in official partners' sites:** As a mean of expanding the dissemination of IOStack, several partners are publicizing the project in their respective websites. For instance, BSC has posted a permanent notice about IOStack⁵. Similarly, MPStor has also integrated publicity of IOStack in its own webpage⁶. Other partners such as URV⁷, Eurecom⁸ and GridPocket have included information on its participation to IOStack project on their corporate website.
- **MPStor Newsletter:** MPStor is also making use of newsletters and mailing lists to spread the word of IOStack across customers and other companies. For instance, the newsletter of February 2015 contains an article about IOStack and it has been pushed to over 4,000 emails⁹.
- **IBM Research blog spot:** IBM has also contributed with a technology blog in which the Storlet framework, as a part of IOStack project, has been discussed. The blog spot can be found at: "Storlets: From research prototype to open source technology"¹⁰ (published April 30th 2015).

¹<http://deim.urv.cat/~dcsm/program.html>

²<http://iostack.eu>

³<https://twitter.com/iostackproject>

⁴<https://github.com/iostackproject>

⁵<http://www.bsc.es/about-bsc/press/bsc-in-the-media/bsc-brings-expertise-data-management-iostack-project>

⁶<http://www.mpstor.com/news-top/news/175-horizon-2020-iostack-project>

⁷<http://ast-deim.urv.cat/web/projects/active-projects/iostack>

⁸<https://distsysgroup.wordpress.com/collaborations/>

⁹<http://www.mpstor.com/media-center-mainmenu/newsletter>

¹⁰<http://ibmresearchnews.blogspot.com.es/2015/04/storlets-from-research-prototype-to.html>

3.3 Industrial Dissemination events

- **DockerCon'15 (Barcelona, November 2015):** Daniele Venzano from Eurecom presented Zoe in a 45 minutes talk at the DockerCon'15 EU Conference held in Barcelona in 2015. The project IOStack was mentioned in the presentation as the environment in which the architectural improvements lead by Zoe will take place. Both the video and the slides of the presentation are available on Youtube and on the Docker Web Site¹¹.
- **OpenStack Israel (June 2015):** The OpenStack Israel event provides diverse tracks on everything OpenStack —from compute, networking, and storage, through hands-on workshops and even a conference that discusses real world stories. In this event, Eran Rom presented “Bringing Compute to Openstack Swift”¹²[3].
- **OpenStack Summit (Vancouver, May 2015):** The OpenStack Summit is a five-day conference for developers, users, and administrators of OpenStack Cloud Software. In this international event, Eran Rom from IBM presented Storlets and their applications to Software-Defined Storage¹³.
- **OpenStack Summit (Tokyo, November'15):** Eran Rom presented “Storlets: Making Swift More Software Defined than Ever” at OpenStack Summit. The talk was about the applications of the Storlet framework to SDS Services¹⁴[4].
- **IBM InterConnect 2015:** IBM InterConnect is a premier cloud & mobile conference, a rendezvous point for users, developers and companies targeted to the cloud business. In this event, Dalit Naor from IBM presented “The perfect match: Apache Spark meets OpenStack Swift”¹⁵. The presentation described the synergies of combining Spark and Swift. In fact, in IOStack IBM is committed to make Swift to cooperate with Data Analytics platforms (e.g., Spark) via active storage strategies for reducing data processing times and boost performance.
- **OneM2M workshop at ETSI (Sophia Antipolis, December 2015):** GridPocket did a talk and demonstration of smart grid data processing solutions at the European Telecommunication Standards Institute¹⁶. The talk was about cloud storage and data privacy with international visitors, which are central points if the GridPocket use case in IOStack.
- **Universal Exposition (Milan, October 2015):** Participation of GridPocket (as a winner H2020 SME Innovation Instrument support) to the side event of Expo Milan, involving talks and demonstration of smart grid data processing, storage and privacy to international companies¹⁷.
- **Cop21 (Lyon, July 2015).** Presentation on energy data privacy and processing by GridPocket at the Lyon EDF pre-event of the Global Climate Summit organized by United Nations.
- **Innovative City (Nice , June 2015).** Demonstration of GridPocket's smart grid data processing solution and presentation of IOStack project plans at the international conference and trade show dedicated to smart cities in Nice¹⁸.
- **MedCop21 (Marseille, June 2015).** GridPocket presentation on smart grid data processing technologies to professionals, politics and institutions of the Mediterranean region countries.

¹¹<http://dockerconeu2015.sched.org/event/88d8ac8cc2193aec4cf73caa7ff76e2b>

¹²<http://www.slideshare.net/openstackil/ibm-swift>

¹³https://openstacksummitmay2015vancouver.sched.org/eran_rom.1sq7eu5t?iframe=no#.VkHrBL90dSA

¹⁴<https://www.openstack.org/summit/tokyo-2015/videos/presentation/storlets-making-swift-more-software-defined-than-ever>

¹⁵<https://www.youtube.com/watch?v=4iez3otZH4o>

¹⁶<http://www.etsi.org/news-events/events/966-2015-12-etsi-m2m-workshop-2015-featuring-onem2m>

¹⁷<http://www.expo2015.org/>

¹⁸<http://www.innovative-city.com/>

- **SmartGrid Paris (Paris, May 2015).** GridPocket technology demonstration, including data processing platform and presentation of IOStack project at SmartGrid Paris professional trade-show¹⁹.
- **Meeting with RTCC (North Carolina, April 2015).** Presentation on European clean tech industry, and research projects, including IOStack at the workshop with Research Triangle CleanTech Cluster.
- **Meeting with SKEMA (Raleigh, April 2015).** Presentation on Smart Grid technologies, and related research projects, including IOStack to students and staff of SKEMA business school in Raleigh.
- **Meeting with ASPROM (Paris, March 2015).** Presentation on BigData, with special focus on data privacy including IOStack project at the (professional conference on Smart networks for electricity, gas and water organized by ASPROM association in Paris.

3.4 External collaborations

In our experience, finding synergies with other projects and companies is an effective way of expanding the impact of an European project. For this reason, our plan during these years is to engage new collaborations and cooperate with other relevant agents in the Software-Defined Storage arena.

At this point, we have few actions in this sense. In 2015, GridPocket has initiated two innovation and research cooperations. On the one hand, in the area of SDS-based privacy and cyber security: “Virtual Private Home” with cooperation with EDF, SMEG, Sagemcom, and CISP. On the other hand, “Collaborative and Confidential Information Sharing and Analysis for Cyber Protection” (H2020 project, starting in 2016) in cooperation of SAP, HP and other research partners. Moreover, GridPocket also presented the IOStack project has during private presentations given to major energy utilities in France, Monaco, Poland, Belgium. The main objective of GridPocket is to publicize the benefits of new SDS technologies on the data processing of new energy meters. Today, there is a big data management and privacy issues that the IOStack project is aiming at solving. The efforts of GridPocket for disseminating our results to other energy utilities companies may potentially attract more use cases as GridPocket to the IOStack agenda.

On the other hand, representatives of BSC have been present in the Network for Sustainable Ultrascale Computing (NESUS)²⁰. The goal of the NESUS Action is to establish an open European research network targeting sustainable solutions for ultrascale computing aiming at cross fertilization among HPC, large scale distributed systems, and big data management. The network will contribute to glue disparate researchers working across different areas and provide a meeting ground for researchers in these separate areas to exchange ideas, to identify synergies, and to pursue common activities in research topics such as sustainable software solutions (applications and system software stack), data management, energy efficiency, and resilience. Representatives of BSC have contacted with other institutions (academics and industrials) in order to disseminate the news of IOStack as well as for finding points in common with other EU projects.

Internally, MPStor has also presented to Intel ISA-L group the IOStack project. MPStor introduced the concept of filtered volumes related to its tasks in IOStack and how ISA-L would be core to the implementation of the filter framework. Via MPStor, we will try to enable collaborations with Intel, which can be of great importance to the impact of the project.

In the following stages of IOStack, establishing collaborations with external companies and EU projects will be a key factor of the dissemination action.

4 Community Involvement and Exploitation Activities

We aim at making IOStack a framework to automate storage provisioning and provide SDS services at large scale. In terms of community involvement, this means that datacenter administrators and

¹⁹<http://www.smart-energies-expo.com/>

²⁰<http://www.nesus.eu/nesus-second-working-groups-meeting-paris-france>

companies with data analytics business processes are our main customer communities. Therefore, our first activity regarding community involvement is focused on testing and deploying IOStack within our use case partners.

On the one hand, we focus on datacenter administrators and operators. In this sense, Arctur is currently working on making the IOStack dashboard suitable and usable for a datacenter administrator. This is a required step towards making the IOStack toolkit an exploitation-ready prototype that could be adopted by other companies. Of course, the role of Arctur is crucial; Arctur will be a pilot customer of IOStack in order to facilitate storage management in the datacenter. The customer experience of Arctur in the following 2 years will lead us to provided a real-world toolkit with actual commercial possibilities. Moreover, Arctur is in contact with other datacenters that may also adopt IOStack, which would be a solid basis for implantation in the market.

Similarly, MPStor is also planning to exploit IOStack in their own business strategy. Today, MPStor is developing automated volume management and a filter framework to deploy advanced SDS services in block volumes. These cutting-edge services represent a value added to their current operation that may lead to advantages w.r.t. existing competitors. To wit, the innovations and developments of IOStack are aligned with the commercial activities of MPStor, which are likely to converge at the end of the project in a success story. In summary, Arctur and MPStor —as well as other potential datacenters— will embody the pilot experience of IOStack in real-world datacenter communities.

On the other hand, we are studying in depth the Big Data necessities of our use-case companies. In particular, Idiada is an automotive company that stores and processes large amounts of information related to vehicle tests. Thus, there are two potential capabilities of IOStack that may enable Idiada of exploiting it as a commercial product. First, the management of large amounts information in Idiada is becoming a problem. To solve such an scalability problem, IOStack promises advanced and software-defined data management techniques, which can reduce the amount of data stored in Idiada's servers. Second, Idiada is also executing parallel simulations of several vehicle crush tests, which requires advanced scheduling and allocation strategies in the compute cluster. IOStack is also focused on providing flexible and intelligent compute cluster strategies via VMs and containers (i.e., Zoe²¹). Currently, we are executing real workloads of Idiada in our experimental IOStack deployments to study the best strategies to solve their problems. In fact, by solving the problems of Idiada in both compute and storage sides, IOStack demonstrates to be a feasible solution to be adopted by many other SMEs around Europe.

As another data analytics use-case, we are working with GridPocket. Essentially, GridPocket works with very large amounts of temporal, comma-separated value (CSV) data that represents the electricity usage of users. In this sense, GridPocket presents two main problems that IOStack aims at solving: i) To efficiently execute complex machine learning algorithms in CSV datasets, and ii) To be able of keeping the privacy and confidentiality of such a sensitive data from external access. The GridPocket use-case opens the door in IOStack to exploit SDS services in cooperation with the compute cluster. This novel perspective of SDS may attract other companies, perhaps outside the energy utility sector, to explore IOStack as a practical tool in their data analytics business processes.

Finally, in the following months of the project we aim at approaching two big open source communities related to the project: OpenStack and Docker communities. Attracting these users will help to disseminate the IOStack project worldwide and we also expect to increase notably the amount of developers and practitioners interested in the project, which may even contribute to the development of the toolkit.

5 Future Dissemination Activities

The immediate future of the dissemination plan of IOStack has two crucial stages: i) Start mass media dissemination, and ii) the OpenStack Summit in Barcelona.

Once the IOStack prototype is ready, we are confident to start doing mass media dissemination and expect users to get interested in the project. Of course, waiting until having a prototype ready is not an arbitrary decision, as timing is very important for mass media dissemination. If such high-

²¹<http://zoe-analytics.eu/>

visibility dissemination arrives too early, the project may be so immature that frustrates the expectations of users and practitioners interested in the project. Conversely, if the mass media dissemination is too late, the novelty of the project may be partially compromised, attracting potentially less users to the project. After our last IOStack meeting in Barcelona, we converged on a integration plan among partners that will be materialized into an integrated IOStack toolkit suitable for mass dissemination and advertising. This stage will be the second year of the project.

From the industrial events perspective, the next DockerCon and, especially, the OpenStack Summit in Barcelona are important targets. Concretely, in the OpenStack Summit of October 2016 we aim at presenting the IOStack SDS architecture for object storage to the community (perhaps other parts of the architecture are also mature enough for being presented as well). This will be a crucial point in which the community will finally know the aim of the project as a whole.

References

- [1] R. Gracia-Tinedo, D. Harnik, D. Naor, D. Sotnikov, S. Toledo, and A. Zuck, "SDGen: mimicking datasets for content generation in storage benchmarks," in USENIX FAST'15, pp. 317–330, 2015.
- [2] R. Gracia-Tinedo, P. García-López, M. Sánchez-Artigas, J. Sampé, Y. Moatti, E. Rom, D. Naor, R. Nou, T. Cortés, W. Oppermann, and P. Michiardi, "Iostack: Software-defined object storage," IEEE Internet Computing, p. To appear, 2016.
- [3] "Openstack (israel, june 2015)." <http://www.slideshare.net/openstackil/ibm-swift>.
- [4] "Openstack (tokyo, november 2015)." <https://www.openstack.org/summit/tokyo-2015/videos/presentation/storlets-making-swift-more-software-defined-than-ever>.