

The eXtreme-DataCloud project - solutions for data management services in distributed e-infrastructures

*Daniele Cesini*¹, *Giacinto Donvito*², *Alessandro Costantini*¹, *Fernando Aguilar Gomez*³, *Doina Cristina Duma*¹, *Patrick Fuhrmann*⁴, *Lukasz Dutka*⁵, *Matthew Viljolen*⁶, *Serena Battaglia*⁷, *Vincent Poireau*⁸, *Luca Dell’Agnello*¹, *Oliver Keeble*⁹, *Rachid Lemrani*¹⁰, *Christian Ohmann*¹¹ and *Jesus Marco de Lucas*³

¹Italian Institute for Nuclear Physics (INFN), INFN-CNAF, Bologna, Italy

²Italian Institute for Nuclear Physics (INFN), INFN-BARI, Bari, Italy

³Istituto de Fisica University of Cantabria (IFCA), Santander, Spain

⁴Stiftung Deutsches Elektronen-Synchrotron (DESY), Hamburg, Germany

⁵AGH University of Science and Technology (AGH-UST), Kracow, Poland

⁶EGI Foundation, Amsterdam, the Netherlands

⁷European Clinical Research Infrastructure Network (ECRIN), Paris, France

⁸Centre National de la Recherche Scientifique (CNRS), IN2P3-LAPP, Annecy, France

⁹European Organization for Nuclear Research (CERN), Genève, Switzerland

¹⁰Centre National de la Recherche Scientifique (CNRS), IN2P3, Villeurbanne, France

¹¹European Clinical Research Infrastructure Network (ECRIN), Dusseldorf, Germany

Abstract. The eXtreme DataCloud (XDC) project is aimed at developing data management services capable to cope with very large data resources allowing the future e-infrastructures to address the needs of the next generation extreme scale scientific experiments. Started in November 2017, XDC is combining the expertise of 8 large European research organisations. The project aims at developing scalable technologies for federating storage resources and managing data in highly distributed computing environments. The project is use case driven with a multidisciplinary approach, addressing requirements from research communities belonging to a wide range of scientific domains: Life Science, Biodiversity, Clinical Research, Astrophysics, High Energy Physics and Photon Science, that represent an indicator in terms of data management needs in Europe and worldwide. The use cases proposed by the different user communities are addressed integrating different data management services ready to manage an increasing volume of data. Different scalability and performance tests have been defined to show that the XDC services can be harmonized in different contexts and complex frameworks like the European Open Science Cloud. The use cases have been used to measure the success of the project and to prove that the developments fulfil the defined needs and satisfy the final users. The present contribution describes the results carried out from the adoption of the XDC solutions and provides a complete overview of the project achievements.

1 Introduction

The eXtreme DataCloud (XDC) project[1] develops scalable technologies for federating storage resources and managing data in highly distributed computing environments. The project is run by a Consortium that brings together technology providers with a proven long-standing experience in software development and large research communities belonging to diverse disciplines: LifeScience, Biodiversity, Clinical Research, Astrophysics, High Energy Physics and Photon Science. The project will end on April 2020 after 30 months of activity. It was funded under the European Union H2020 framework program (Call EINFRA-21-2017 Research and Innovation action) with a total EU contribution of 3.07 million Euros. XDC developments enrich already existing services for data management in e-infrastructures adding new functionalities as requested by the research communities represented in the consortium and targeting the computing infrastructure built by European Open Science Cloud (EOSC) [2], the EGI Foundation (EGI) [3], the Worldwide LHC Computing Grid (WLCG). An introduction to the project, the description of its objectives, the involved research communities, the services composing its software catalogue and the foreseen technical architecture can be found in [4][5]. The next section provides a brief project overview, while Section 3 reports the main development achievements of the projects.

2 The XDC project overview

XDC is a follow up, in the field of data management, of the INDIGO-DataCloud project[6] and continues the INDIGO work on storage Quality of Services definition/handling and on data lifecycle management. It develops smart orchestration tools to realize a policy driven data management in heterogeneous e-infrastructures. XDC develops the building blocks of the software layer that can be used to implement the so-called Data-lakes, intended as storage federation providing transparent or quasi-transparent access to data stored in geographically distributed huge sites connected through a high-bandwidth network backbone. XDC is a user driven project and table1 lists the communities represented in its Consortium.

Table 1 - XDC represented User Communities

Community	Domain	Reference partner	Role in the project
ECRIN: the European Clinical Research Infrastructure Network	Medical research and investigation of clinical trials	ECRIN ERIC [7]	Driver for the developments on metadata managements tools, easy to use and web based, that allow researchers to aggregate centrally metadata coming from heterogeneous sources
WLCG: The Worldwide LHC Computing Grid [8]	High Energy Physics	European Council for Nuclear Research (CERN) and Italian Institute for Nuclear Physics (INFN)	Requirements for the implementation of “DataLakes” to simplify the dynamic extension of sites to remote locations to allow transparent remote datasets access through smart caching mechanisms. QoS support request for the sites composing the DataLake.
XFEL: the European x-ray free-electron	Photon physics	German Electron	Driver for the storage QoS management tools and for the storage

laser facility [9]		Synchrotron (DESY)	events notification system needed to process data coming at high rates.
Lifewatch: e-Science and Technology European Infrastructure for Biodiversity and Ecosystem Research [10]	Biodiversity	University of Cantabria (UC)	Requires the integration of different tools to manage the data lifecycle, the production of data based on FAIR (+R, Reproducibility) principles and the automation of storing in distributed environments data and metadata produced by sensors of different sources (satellites, simulation models, air/water quality monitoring etc)
CTA - The Cherenkov Telescope Array [11]	Astrophysics and Astroparticle physics	French National Centre for Scientific Research (CNRS)	Requires the creation of a distributed storage infrastructure to store, with multiple QoS, data coming from the two CTA antennas. Metadata extraction and manipulation is also need during the process.

The XDC “toolbox” used to address the requirements provided by the User Communities is a set of services already available as well-know production quality tools or prototypes existing at least at TRL6 (Technology Readiness Level [12]). All XDC products are released with open source licenses at least at TRL8. Table2 summarizes the tools used in the XDC architecture and developed with the contribution of XDC, which developments aimed at improving the functionalities, the performances, the usability and scalability of the all the listed components.

Table 2 – The XDC « toolbox » - the services developed by XDC partners and/or used to implement the XDC architecture for data management.

Tool/service	Main developer in XDC	Functionality provided and role in the XDC architecture
Onedata[13]	AGH University of Science and Technology in Krakow	Virtual filesystem using data resources backed by datacentres and storage providers worldwide. Provides metadata management capabilities.
dCache[14]	German Electron Synchrotron (DESY)	Storage backend system: a Data Storage and Management system developed by dCache.org deployed in more than 60 sites around the World, including 8 WLCG Tier 1 centres.
EOS[15]	European Council for Nuclear Research (CERN)	Storage backend system: scalable to many tens of petabytes is managing the data of the distributed WLCG Tier 0 centre.
StoRM[16]	Italian Institute for Nuclear Physics (INFN)	Storage backend system: provides SRM interface to shared filesystems like SpecrumScale (formerly GPFS) or Lustre.
Rucio[17]	Initially developed by the ATLAS collaboration at CERN	Manage accounts, files, datasets and distributed storage systems - is the data management policy engine for XDC.
INDIGO PaaS Orchestrator[18]	INFN	Allows to instantiate resources on Cloud Management Frameworks at the PaaS level through TOSCA YAML Profile 1.0 [19]. Used to steer data movements based on policy via Rucio and FTS.
Dynafed[20]: The Dynamic	CERN	Provides a very fast dynamic namespace that it exposes via HTTP and WebDAV. Clients can

Federations		browse the Dynamic Federation as if it were a unique partially cached namespace
Xrootd proxy cache ("Xcache")[21]	No development, deployment recipes (INFN)	Caching technology for the xrootd protocol
NGINX[22]	Plugin development (INFN)	Caching technology for the httpd protocol
INDIGO-IAM[23]	No development, usage (All partners)	Identity and Access Management system based on OpendIDConnect originally developed by INFN within the Indigo-DataCloud project

3 Project Achievements

After an initial phase of requirement gathering which involved all the scientific communities represented in XDC, the project produced two major releases: XDC-1, codenamed Pulsar, released in January 2018 and XDC-2, codenamed Quasar, released in March 2020. Both releases are based on the general architecture that is presented in [4] and improve or add missing functionalities to the set of software components reported in Table2, addressing important topics like federation of storage resources, smart caching solutions, policy driven data management based on Quality of Service, data lifecycle management, metadata handling and manipulation, optimized data management based on storage events. The two XDC releases allow a tighter integration between the participating components in order to enable workflows which exploit the whole data platform as a coherent infrastructure rather than a set of disparate services. Table3 provide a not comprehensive list, component by component, of the main new functionalities and enhancements provided by XDC. More details and an all-inclusive list can be found in the project releases documentation available at [24].

Table 3 - Main developments and new functionalities added by XDC to the released components

Component	XDC Achievements
Onedata	Performance and scalability improved: 100M+ file per collection, 100GB/s local processing throughput, 100Gbs distributed mesh transfer, Simplified deployment with Onedatify, New GUI interface, advanced group/role right management, WebDav driver for dCache integration, internal map/reduce system for large scale indices. ECRIN plugin to inject data from various medical sources, metadata management subsystem improved performances, integration with external databases for storing metadata. OpendIDConnect (OIDC) support.
dCache	Storage events support with Kafka and SSE, support to notify events, new plugin for SSE, clients can discover changes in dCache namespace using an interface modelled after the inotify API. dCache View is updated, 3rd party copying functionality more robust and scalable.
Dynafed	OIDC support, both as a Relying Party (redirecting a browser to an IdP) and Protected Resource (consuming oauth access tokens for non-interactive access), facilitating in this way the integration with the XDC Orchestrator and allowing browser based access without X509 certificates. This is configured through Apache's mod_auth_openidc. "Fourth party copy" - dynafed can function as the active party for data distribution. This enables services without third party copy support (such as S3) to participate fully in the data distribution infrastructure
EOS	Addition of QoS classes and QoS API when interacting with namespace entries, CDMI gateway for QoS transitions, Introduction of the Converter Engine, the EOS component responsible for scheduling and performing file conversion jobs. A conversion job means rewriting a file with a different storage parameter: layout, replica number, space or placement policy
FTS	Added OID connect support, on the basis of OAuth2 tokens issued by an

	Authorization Server such as IAM. full support of managed QoS transitions.
RUCIO	Authentication and authorization mechanism extended to support (JWT) tokens using OIDC protocol, to allow permission control downstream (Rucio → FTS3). Implemented internal mechanism using token exchange and token refresh grants. Added user pre-provisioning (via new Rucio SCIM client) implemented as a ‘Rucio probe’ script. New Rucio daemon implemented taking care of token deletion, token refresh and clean-up of expired authentication OIDC sessions. Third party copy supported in the entire chain Rucio → FTS3 → dCache
Caching-on-demand	Provides recipes and PaaS description templates for an end to end deployment of an XCache cluster. Container based, support for Docker, Kubernetes, Ansible recipes.
PaaS Orchestrator	Added OID Connect support in the interaction with Onedata and Rucio, description Dynafed resources added in TOSCA templates, added scheduling strategy using Dynafed, implemented Cloud Providers Retry Logic, improved description of Onedata spaces, Added the possibility to steer data movement using Rucio. Improved web interface to submit and monitor orchestration requests: “Paas Orchestrator Dashbord”.
TOSCA Types and Templates	Extended TOSCA Simple Profile in YAML Version 1.0 to add high level entities. Added new types for Onedata and Dynafed storage resources, added new types for describing a Kubernetes cluster. New Ansible roles implemented, new type for describing a JupyterHub node, updated example templates for deploying Chronos dockerized jobs that use Onedata for managing input/output data.

Conclusion

XDC focused on the continuous interoperability and highest scalability of a well-established toolbox to build a European wide data storage and management system for science. Successful efforts have been undertaken to involve community and industry standardization organizations (i.e. RDA) as well as large, well established, data driven communities to guarantee sustainability of the XDC achievements. Scalability is pursued not only through individual components but also through their interactions, so the XDC architecture has been implemented making its components interoperable, for what concern functionalities but also modern and fine grained authentication methods, allowing a distributed model to be fully exploited and the horizontal scaling at the infrastructure level achieved. XDC is active in all the forums and boards where the represented communities discuss the requirements and the implementation of the used e-infrastructures presenting the solutions developed in the last 30 months. Moreover, to guarantee sustainability, all new features have been pushed to their upstream repositories, leading to a wide deployment as soon as the new versions are downloaded and installed from those sources. XDC is also interacting with the projects and entities that are shaping the European Open Science Cloud (EOSC), promoting the developed solutions as the building blocks of the data management systems of this emerging infrastructure.

References

1. XDC project, “XDC” [web page], www.extreme-datacloud.eu [accessed 2020-03-13]
2. EC EOSC, “EOSC” [web page], <https://ec.europa.eu/research/openscience/index.cfm?pg=open-science-cloud>, [accessed 2020-03-13]
3. EGI Foundation, “EGI” [web page], <https://www.egi.eu/>, [accessed 2020-03-13]

4. D. Cesini et al, “Advancements in data management services for distributed e-infrastructures: the eXtreme-DataCloud project”, *Proceedings of 23rd International Conference on Computing in High Energy and Nuclear Physics (CHEP 2018)*, EPJ Web Conf. 214 04044 (2019), doi: 10.1051/epjconf/201921404044
5. D. Cesini et al, “The eXtreme-DataCloud project: data management services for the next generation distributed e-infrastructures”, *Proceedings of Conference Grid, Cloud & High Performance Computing in Science (ROLCG)*, **1–4** (2018) doi: 10.1109/ROLCG.2018.8572025
6. D. Salomoni et al., “INDIGO-DataCloud: a Platform to Facilitate Seamless Access to E-Infrastructures,” in *Journal of Grid Computing*, vol. **16**, number 3, Sept. 2018, pp. 381–408
7. ECRIN ERIC, “ECRIN” [web site], <https://www.ecrin.org>, [accessed 2020-03-13]
8. WLCG infrastructure, “WLCG” [web page], wlcg.web.cern.ch/, [accessed 2020-03-13]
9. XFEL Organization, “XFEL” [web site], <https://www.xfel.eu>, [accessed 2020-03-13]
10. Lifewatch ERIC, “LifeWatch” [web site], <https://www.lifewatch.eu>, [accessed 2020-03-13]
11. CTA Experiment, “CTA” [web site], <https://www.cta-observatory.org>, [accessed 2020-03-13]
12. TRL definition in the H2020 General Grant Agreement, “TRL”, [web site], http://ec.europa.eu/research/participants/data/ref/h2020/other/wp/2016-2017/annexes/h2020-wp1617-annex-ga_en.pdf, [accessed 2020-03-13]
13. onedata.org, “Onedata” [software], version v18.02.0-rc13, 2019. Available from <https://github.com/indigo-dc/onedata>, [accessed 2020-03-13]
14. dCache storage, “dCache” [software], version vXDC-1,2019. Available from <https://github.com/indigo-dc/dcache/tree/xdc>, [accessed 2020-03-13]
15. EOS storage, “EOS” [software], version v4.4, 2019. Available from <https://github.com/indigo-dc/eos>, [accessed 2020-03-13]
16. StoRM storage, “StoRM” [website] <http://italiangrid.github.io/storm/index.html> [accessed 2020-03-13]
17. Barisits et al., “The ATLAS Data Management System Rucio: Supporting LHC Run-2 and beyond”, ACAT, Seattle, 2017
18. INDIGO Paas Orchestrator, “Orchestrator” [software], version 2.1.1-FINAL, 2019. Available from <https://github.com/indigo-dc/orchestrator>, [accessed 2020-03-13]
19. YAML Profile 1.0 specification, “Tosca” [website], <http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.0/cspr01/TOSCA-Simple-Profile-YAML-v1.0-cspr01.html>, [accessed 2020-03-13]
20. Dynafed federator, “DYNAFED” [software], version v1.3.3, 2019. Available from lcgdm.web.cern.ch/dynafed-dynamic-federation-project, [accessed 2020-03-13]
21. XCache service, “XCache”, [website], <http://slateci.io/XCache/>, [accessed 2020-03-13]
22. NGINX web server, “NGINX” [web site], <https://www.nginx.com/>, [accessed 2020-03-13]
23. A. Ceccanti et al. “The INDIGO-Datacloud Authentication and Authorization Infrastructure”, *IOP Conf. Series: Journal of Physics: Conf. Series* **898** (2017) 102016 doi:10.1088/1742-6596/898/10/102016
24. eXtreme-DataCloud project release documentation, “XDC Documentation” [website], <https://releases.extreme-datacloud.eu/en/latest/index.html>, [accessed 2020-03-13]