

DPM Community White Paper

The Disk Pool Manager community ([list of signers](#) at the end of the document)

Introduction / Abstract

The idea of a DPM community white paper was put forward during the 2019 DPM Workshop in Bern [1,2]. The triggering event was the announced reduction of full-time equivalents (FTEs) in the CERN DPM development group. This cast some doubt on the sustainability of the DPM project itself, particularly since this disengagement happens while sites, on the contrary, are engaged in various DOMA related R&D activities [3,4] and still planning important investments on their DPM infrastructure. There is therefore a general request from sites for a common long-term vision which is shared within the overall DPM collaboration.

The purpose of this document is to provide CERN, WLCG, EGI and the funding agencies with collected feedback from DPM sites, specifically their vision and their plans for the future as well as some statements about the project sustainability from the community point of view.

It should be stressed that this document concerns an important amount of storage resources. About 100PB, roughly 40% of WLCG Tier-2's resources are today running on DPM storage systems, mostly hosted in small/medium size Tier-2's.

DPM is popular because it provides stable and performant grid storage, while keeping administration costs quite low. Moreover, the developers provide extremely reactive support and they have been capable of creating, over the years, a real community by encouraging exchange among the administrators. Both these points are of fundamental importance for small/medium sites with only one or two administrators, who are often quite overloaded.

It should be also pointed out that, by allowing for the deployment of grid services with a reasonably low initial investment (in particular in terms of manpower), DPM brings to WLCG (and ultimately to CERN) some important benefits. In particular, it helps to incorporate resources from smaller national or regional communities, de facto favoring the exploitation of additional "local" budgets.

Moreover the DPM team has recently released a new DPM software core - called DOME (Disk Operations Management Engine) [5] - which dramatically improves performance and stability. With DOME, sites can reasonably expect to scale up to 10-20PB without major problems and to ~100PB with some more effort. On this basis most of the DPM sites are confident on the possibility to keep running DPM at least for the next 5 years - while roughly doubling their size - without major issues. But, this of course also depends on the availability

of a reasonable number of FTE in the development team to assure high level support and urgent developments.

Longer term plans are more difficult to assess as they depend critically on the definition of the WLCG data management in the HL-LHC era. Medium size Tier-2 storage may simply no longer have a role in the WLCG architecture, or it will come with a number of new requirements. Whether the DPM development team will be capable of supporting the software evolution needed by these new requirements is one of the great unknowns and major concerns for the DPM community.

One last point that should be stressed is that a massive migration of DPM sites to a different storage technology - given that a suitable product is available - would be a considerable cost for sites, mostly in terms of manpower, and will necessarily translate in less resources (storage) available for WLCG.

So far, the DPM community has made an important investment and has put considerable effort into the project and into contributing to the collaboration itself. By writing this document, DPM sites hope to trigger some discussions among all those concerned, in order to successfully plan the future steps of the evolution of the storage infrastructure. Some feedback is expected from the other involved communities (CERN, EGI, WLCG, experiments and funding agencies) about their views and their plans. In particular, DPM sites, in order to assess their medium and long-term plans, need to understand to what extent CERN will maintain and reinforce DPM support and development in future years, as this critically impacts the possible scenarios.

Current Status

The Worldwide Infrastructure


LCG-DM/DPM has been part of the gLite/EGI middleware for more than 10 years. Over this time it has been by far the preferred storage technology for small and medium sized Tier-2 sites and it can still be considered the most reasonable solution for sites looking for an easy to use and stable, while still performant and scalable, storage solution.

There are currently about 85 DPM sites, in more than 25 countries, providing roughly 100PB of storage (for comparison the total LHC Tier-2 pledged disk is around 240PB). Most of these sites are part of the WLCG infrastructure and contribute with an important fraction of the overall grid storage available to LHC experiments. Besides LCG, there are more than 40 grid sites in 15 countries (mostly in Europe but also outside) which are using DPM to support the storage-related needs of various other EGI communities.

National communities mostly contributing in terms of storage capacity are: UK (12 instances, 24PB), France (14 instances, 20PB), Japan (1 instance, 10PB), Taiwan (5 instances, 10PB), Italy (12 instances, 7PB), see Fig.1.

While a big part of this community is still composed of “small” sites, there are 35 instances above 1PB summing up to a total of 90PB. The largest single DPM instance to date is deployed by the University of Tokyo and it provides 10 PB of disk [6]. Over the last few years the general trend has been to consolidate toward fewer sites with bigger - multi PB - storages. The future of DPM is probably in the medium-large sites which are currently above 2PB, like the WLCG Tier-2 site in Prague now providing 4PB.

Fig 1: DPM storage provided by different national communities.



Most of deployed infrastructures are based on rather standard hardware: multiple disk-servers each with one or more Raid-6 raid-sets. Disk sizes range from 3TB to 14TB - depending on server generations - and connection bandwidth from 1Gbps to 40Gbps per server, with 10Gbps servers being the most common case.

There is one important point about this community of sites which should be stressed. DPM sites contribute to 40% of WLCG Tier-2 storage resources. They are in some cases funded by national institutions that cannot afford bigger - Tier-1 size - grid centers. In other cases they belong to national communities that already have a big national center but also have a cloud of regional centers co-funded (sometimes up to 50-70%) by local institutions. Thus, in most cases, they represent resources that could only be attracted by integrating into WLCG small and medium size centers (up to a few PBs) with limited infrastructure and manpower. The DPM project - a grid storage technology with production level performance, low operational costs and a very active support - had and still has a key role in making this possible.

The Storage Technology

Over the years, DPM has proved to be a stable and well performing storage solution. Thanks to the outstanding work done by the DPM development team at CERN, sites deploying DPM have been able to provide, to their supported VOs, storage services with a high level of reliability and a complete integration in WLCG/EGI operations.

The reengineering done in the past 3 years by the developers has been successful in delivering the new DOME software stack. This new implementation provides all legacy-DPM functionalities except the SRM and RFIO services, and meets much higher performance and stability standards. DOME allows for fully SRM-less storage operations; the SRM interface is officially no longer supported since October 2019.

The new stack was extensively tested and it is already deployed in many production sites, while the remaining sites have been ticketed to get them to upgrade as soon as possible. In all cases it provided higher stability and increased dramatically the performance for all SRM-less operations.

A concrete example of how DOME enhances DPM scaling in terms of performance is provided by the Tokyo Atlas site. Running a single 10PB instance, Tokyo explored the limits of the legacy DPM implementations and has needed to invest a lot of effort to maintain adequate SRM performance. Moving to DOME and switching to protocols other than SRM has radically improved the situation and proved that we can confidently scale up to DPM sites of several tens of PBs. Along the same lines, intensive stress tests performed at Prague Tier-2 showed that metadata operations are at least one order of magnitude faster with DOME than with SRM [7]. The load on the database by SRM metadata operations is just one of the bottlenecks that limit the performance of the legacy stack and DOME removes these limitations. Such results support the developer team statement that DPM sites, after migration to DOME, can safely evolve and grow one order of magnitude before meeting scalability problems.

Still, the new stack misses some features which would be suitable for sites wanting to evolve to larger storage endpoints: head-node and/or database high availability or horizontal scalability, for example, as well as support for JBOD and object storage (the current “standard” Raid-6 based HW configuration is approaching its limits). However the plugin based architecture of the codebase already has the flexibility needed to provide such features. For example, a plugin for S3 based backend, has been used for a few years, in particular by BELLE II dynamic data federation based on Dynafed [8].

Besides the storage system itself the CERN development team also provides Puppet configuration modules. These are released with the core system packages. These modules are extremely important for sites since they considerably simplify the deployment and configuration of the storage system as well as the inclusion of the new functionalities.

The Collaboration

The DPM collaboration [9] started in 2013 with the emergence of the DMLite plugin-based framework announcing the complete reengineering of the LCG-DM legacy stack (e.g. rfio,

libshift, SRM...) into DOME (Disk Operations Management Engine) which has been achieved in 2017.

The purpose of the DPM collaboration was to ensure the maintenance, development and support of the DPM software stack for the benefit of the user communities and the sites. From the beginning, it was agreed among the partners (e.g. countries, institutes, projects, sites or individuals) that development activities will cover software maintenance and feature implementation. It is perhaps useful to recall that the development of the new DOME stack was motivated by the needs for DPM stability and long-term support. Beyond that goal, the collaboration was expected to attract new contributions and to find the necessary effort for possible new large scale development projects. It is fair to say that this goal was not met.

During the past six years, the CERN DPM team has undertaken all the critical tasks including development, certification, integration, release and all necessary maintenance. CERN has also contributed to support, documentation and configuration, and provided support for all collaborative development. CERN, Czech Republic, France, Italy and the UK have remained engaged as key partners. It is also fair to say that the vast majority of contributions outside CERN came from site admins doing deployment on testbeds and testing of the new DPM versions e.g. right after release in EPEL-testing repositories. It is important to note that these contributions have been made on a “best effort” basis :

- UK has developed administration tools, provided Ansible configuration and also worked on Argus integration.
- Italy has contributed to the EPEL-testing activity on testbeds and advanced feature validation (multi-site installation, volatile pools, DynaFed as HTTP Dynamic Federation).
- France has also been involved in EPEL-testing and also Quattor/Puppet configuration setup and ATLAS and CMS validation.
- Czech Republic (Prague Tier-2) as an early DOME adopter has debugged and fixed issues related to the deployment in production, and also put effort into improving scripts used for the upgrade to DOME and easier integration with WLCG/EGI.

It is worth noting that EGI has not been explicitly part of the DPM collaboration but has been testing and distributing DPM as part of its middleware distribution. EGI is providing some infrastructure services useful to the DPM community and users such as the GOCDB, UMD repositories, ARGO, APEL and GGUS. EGI has also been doing coordination of updates across the infrastructure as well as supporting and interacting with user communities.

Enabling DOME, supporting CentOS 7 and Puppet configuration and making a legacy-free, SRM-free DPM have been the main goals of the DPM collaboration since 2016. The development team at CERN has secured DPM as an efficient storage solution which is able to cope with the upcoming technical and operational challenges (e.g. higher load, scaling up, TPC, WebDAV, macaroons, XRootD, multi-site, caches, etc.). Furthermore, an intensive and collaborative testing activity was done by sites and the tremendous amount of exchange on the DPM forum has played in favor of a healthy DPM community.

Nevertheless, there is no certainty that the DPM collaboration is, as it is, capable to address the long-term challenges and there is a need to estimate the requirements in development and support and to identify a sustainability mode for the future of DPM. The possibility of a manpower decrease at CERN by the end of 2019 is one of the concerns that should be taken into consideration. In case a long-term road-map for DPM would be defined, there could be opportunities for additional contributions to the DPM collaboration coming from contractual positions of ongoing or future EU projects.

Beside the DPM collaboration, the WLCG DPM Upgrade Task Force [10] has been started in 2018 to coordinate the migration process to DOME including upgrade to newest DPM version, reconfiguration to enable DOME and SRR (Storage Resource Reporting) support. There are clear benefits in setting up such an ad-hoc task force within WLCG and this working mode may help in the future as well.

Statements for the Future

Sites evolution

Being one of the most popular solutions for small/medium size storage, DPM has so far played an important role in maintaining technical coherence in the population of WLCG Tier-2's - at least in Europe and Asia. The approaching structural limits of some technologies (e.g. raid-6), the evolution of the WLCG model (e.g. DataLakes) and doubts on the sustainability of the DPM project itself may spoil this coherence in the upcoming years.

Following the trend of the last few years, and in line with the direction indicated by LHC experiments, the DPM sites community is progressively evolving into a smaller number of relatively bigger centers. Most of the medium size sites are planning to roughly double their capacity in 5 years. Other groups of sites are planning to consolidate multiple instances into a single distributed storage endpoint. In almost all cases, these sites foresee to evolve while keeping DPM as their storage technology. It is thus likely that in 5 years from now the community will count fewer DPM instances, each spanning from 5 to 20 PBs (as mentioned above, we expect DPM to be stable and well performing at these scales).

If this is the general trend, there are some noticeable exceptions, for example UK sites. Small UK sites are planning - encouraged by experiments - to move to cache only storage, for example using XCache. At the same time, some of the largest UK Tier-2s are concerned by the sustainability of the standard raid-6 backend solution - which is so far the only one available in DPM - and are evaluating the possibility to migrate to industry-standard storage solutions which support distributed resilience, such as, for example, Ceph.

Far less can be said on the evolution on a longer scale. Much depends on the outcome of the DOMA project, which is to define the WLCG data management architecture for the HL-LHC era.

Medium size disk storages may no longer be part of the WLCG model (but local communities and end-users will still need dedicated disk space). On the contrary, if there will still be a role for this type of resources, this will come with some new requirements in order

to be integrated into the architecture. Developing such new functionalities will require manpower and the DPM development team may or may not be in the condition to cope with this.

It should be stressed that the scenario in which DPM sites will be forced to migrate to new technology for lack of development manpower in DPM will bring important consequences. The net effect of this will be less storage resources available globally (recall that DPM sites sum up to 40% of WLCG T2 resources) as some sites may just renounce providing WLCG storage while others will have to divert a part of their financial resources to the migration effort.

It is also worth mentioning that the lifetime of the Hardware (HW) - approx. 7 years - puts some constraints on today's investments in order to match possible technology changes in the future. This is particularly true for Tier-2s that work with yearly (flat) budget, progressively renewing their HW and increasing their capacity. Other sites have bigger point-wise investments - for example every 3-4 years - and may more easily make disruptive technical choices. Which of these two patterns applies to a given grid site is however determined by the policy of its funding agency and it is, in general, non-negotiable. Thus, when considering the plans of sites up to the HL-LHC horizon we have to take into account that some of them may still postpone technological choices for a few years while others already need to think about the coherence between their current purchases and their future evolution.

Another point is that the storage capacity will probably grow but the Tier-2 manpower is unlikely to increase. Consolidated solutions that reduce operational costs are thus to be preferred. **So far DPM has indisputably been one of such solutions, whether this will or not be the case in the future is part of the community concerns about the project sustainability.**

Medium-term view of DPM evolution and R&D activities

For the next five years, some of the upcoming technical evolutions are known and some operational concerns are expressed by medium size sites which are expecting less expensive managed and consolidated storage. Beyond that, the site evolution and R&D activities will also drive medium term requirements.

Benefiting from the DOME software, some of the medium-sized sites intend to continue operating their DPM storage as it is (i.e. non-SRM and xrootd/http only), while others may consolidate at a regional or national level, different DPM endpoints into a single federated instance with a single namespace and distributed data pools. Some DPM sites are also part of R&D activities, gaining experience with distributed storage or investigating possible use cases of the integrated DPM caching mechanism.

R&D activities involving DPM and DPM sites: Some DPM sites are already involved in ongoing activities related to data access in the contexts of WLCG DOMA as well as ESCAPE [11,12].

- First of all, it is important to recall that DPM is part of DOMA ThirdPartyCopy (TPC) and QoS activities [13,14], with important contributions from both the CERN DPM and Prague teams.
- In the UK, there is an evident willingness to continue working on the central banning with DPM and already progress in the Argus integration.
- In Italy at INFN sites, DPM volatile pools and distributed configuration, also in conjunction with Dynafed federation, have been tested in the environment of different experiments like ATLAS and BELLE II [15,16]. DPM built on a Ceph file system is also under testing. Such activities will continue in the Data Lake context and in particular within the ESCAPE project.
- In France, DPM sites are part of the IN2P3 project contribution to DOMA, called DOMA-FR, and evaluating performance and cost of remote access and future distributed storage services. A DPM federation (FR-ALPES) has already been implemented in the south-east region and integrated within the ATLAS environment [17]. The objective is to quantify the necessity to use caching mechanisms to optimise the data access between close sites. Other key points to be considered are storage consolidation, scale-out, sharing of critical services and reduction in operational complexity and cost. This infrastructure is foreseen to be a testbed for a DPM component within the ESCAPE Data Lake prototype. The goal of these R&D activities is to identify the potential limitations or possibly further needed developments of the DPM software. From a different angle, it is foreseen to evaluate Ceph as an alternative technology interfaced to the Grid through bridges supporting current protocols (XRootD,...). The current model is the ECHO project [18] at RAL but, if possible, with the redundancy implemented across grid sites.

In any case, various DPM sites are willing to take part in the integration activities of new components and will probably be gaining experience with the new token-based WLCG AAI quite soon. In that respect, the support for token-based access has been identified as the major short term requirement :

- **Support for token-based access:** According to current WLCG and OSG GSI migration plans [19], support for GSI can start to disappear pretty soon and unless there is a token to X.509 translation service, storage implementations should add support for WLCG/SciTokens. Non X.509 based access is also a ‘must-have’ nowadays for solutions proposed by EGI, especially those incorporating a distributed data store, and such support within DPM is more than desirable (together with WebDAV and TPC).

Token-based access can become a critical requirement for storage to be able to participate in a distributed data store. DPM already supports XRootD and WebDAV-based third-party transfers. Experimental support of macaroons has been enabled as the first example of “bearer tokens”, to provide a solution for delegated authorization. A first experimental version supporting token-based authentication may also be available soon. So DPM is rather in a

good position, but full support for the WLCG future token-based authorization will require some work in adaptation and integration and testing.

- **Central banning / Argus integration:** Centralized banning still depends on legacy DPM code. Argus integration with DOME flavour of DPM has recently raised discussions. Moreover the Argus integration should evolve with the Token support.

Another point that was raised by a number of sites is related to DPM scale-out and high-availability capabilities for the next 4 or 5 years.

- **DPM scale-out and high-availability capabilities:** The need for horizontal scalability of the head node has been raised and identified as a potential requirement for the future. As noted above, the sites are quite confident with the performance of metadata operations. However, while the DPM database can be configured with replication, the main concern with big distributed pools remains that the DB becomes huge and quickly unmanageable. One possibility could be to support a kind of “multi-head DPM” installation with perhaps different databases taking care of different subtrees of the namespace (e.g. subtrees of the different VOs basically never interact...). Running the DPM head node in a High Availability configuration would be a clear benefit, also for the DPM maintenance. But the possibility of configuring a DPM head node “active/passive” cluster has never been tried.

In conclusion, sites would be interested to discuss this topic within the DPM community and to know which way High Availability configuration makes sense from the developers’ point of view.

Long-term view of DPM evolution

Long-term evolution plans of DPM sites depend, of course, on which will be the actual WLCG architecture at the HL-LHC time scale. The DOME project is currently working on the definition of the Data Management for HL-LHC and has not yet delivered any definitive statements. Therefore there are several questions about the long-term future that still cannot be addressed: will the medium size Tier-2s still participate in the WLCG storage? If so, in which role? Which functionalities will be required for a site or storage endpoint to be part of the WLCG “Data Lakes”?

Once the answers to these questions are known, the DPM sites will be able to evaluate whether running a medium size grid storage is still a fitting choice in the context of WLCG and if DPM will be able to evolve and provide all the required functionalities. The latter critically depends on the development effort available.

On the other hand, despite these uncertainties, some general lines of evolution can already be identified:

Caching: Sites with cache storage are very likely to be part of the future data lake architecture. DPM currently offers a feature called Volatile Pool, providing a cache-like behaviour driven by custom scripts. This full file caching mechanism released in 2017 will

probably have to be improved in order to be used in a production context. For example more powerful and customizable cache algorithms may be implemented as well as read-ahead.

DPM over different backend solutions with different resilience profiles: Currently DPM implements a very simple file based distributed file system. Data resilience is delegated to the underlying storage which in all production instances is a set of RAID6 disk servers. All reasonable projections of future storage resilience suggest that RAID6 will be insufficient to guarantee data resilience on a single server within a matter of few years. A natural evolution would be to run DPM over a backend implementing an actual distributed file system, for example Ceph. This would allow distributed resilience and also increase the performance through parallel access over multiple servers.

Related to this is the possibility to run in Just a Bunch of Drives (JBOD) mode, removing all resilience and letting the experiment workflows manage data loss gracefully. This is very closely related to the implementation of cache storage.

Support for object storage: some of the forthcoming HEP experiments are currently looking at the advantages/disadvantages of object based data management solutions. Supporting this type of storage as a possible backend will probably require some important development work. It would be worth estimating the status of the existing S3 backend plugin for DMLite and the amount of work required to support S3-like technologies in future.

Quality of Service (QoS): among the possible outcomes of DOMA will be a set of QoS classes that would allow the WLCG workflows to classify storage services in a more complete way (currently we only distinguish between “tape” and “disk”). Small and medium sites will probably be less concerned by this. Still it is not excluded that also disk-only WLCG storage systems will be required to provide a given number of QoS classes.

As already noted, several DPM sites are currently participating in DOMA related R&D activities.

Statements on DPM sustainability

Some conclusions can be drawn from what has been detailed in the document. Here are some final statements about sustainability in the near and medium term future of the DPM project from the point of view of the Tier-2 sites currently having this technology deployed.

DPM as storage technology:

- Thanks to the recent releases of the legacy-free and SRM-free new flavour of DPM (DOME), there is a broad consensus on the fact that **DPM is now able to meet much higher performance and stability and should remain an appropriate storage technology for Tier 2 sites at least for the next 5 years** unless major unknown changes arise.
- However, the need for new developments will arise as soon as new requirements come from operations. For example, in the short term, DPM sites express the need for token-based access in order to support future WLCG authorization methods. **The presence of an active development team with sufficient manpower to timely**

provide the required software evolution is thus a mandatory condition for the sustainability of the DPM infrastructure.

Future of DPM is in the medium-large sites:

- Most of the medium size DPM sites intend to continue to invest in their DPM infrastructures in the coming years. Their plans are to roughly double their capacity on a 5 year term but also to consolidate multiple instances into a single distributed storage endpoint.
- There is a tendency for the smaller sites providing less than 1 PB to move to a federation or to cache-only storage.

Alternatives to DPM, potential cost and impact of migration:

- DPM sites are satisfied with DOME/DPM performances and happy with the organisation and responsiveness of the DPM support. They have a longstanding experience in managing the DPM software and they have just invested effort in the DOME migration. For these reasons the **great majority of DPM sites are not in favour of a migration to a different storage technology in the short-medium term.**
- If however needed, several alternatives to DPM are available (dCache, XRootD, StoRM, EOS, Echo). So far there is no favoured option in the community. In any case the DPM sites are convinced that **the cost - in terms of money, manpower and time - of such a migration should be considered with great attention and that the impact assessment should probably include the experiments, other VOs and funding agencies.**
- Some sites working with annual flat budget have serious concerns regarding a change of technology which may require different types of hardware and/or a peak of investment in order to run both storage infrastructures at the same time. **Providing storage services requires planning over 5-6 years and thus a clear visibility, on the same time scale, of the future of the technology deployed. Having such visibility on DPM future is thus a necessary condition for the sustainability of the supported storage infrastructure.**

R&D activities:

- Italy, Czechia, France,... and the UK have ongoing R&D activities involving DPM technology or planned contributions to DOMA and/or ESCAPE, which is a very positive sign.
- Contribution to DPM R&D could be of interest for contractual positions of ongoing or future EU projects. But this would require to suitably define priorities and long-term road-map for DPM .

DPM collaboration & future dev. contributions:

- During the past six years, the CERN DPM team has undertaken most of the work to provide the new software stack based on DOME. On the other hand, recently CERN has announced a reduction of the DPM team FTEs by the end of 2019. Now, **DOME is a major achievement which may be converted into Petabytes of performing storage for WLCG in the medium term but this requires a clear view of the evolution of the DPM support and development on the same time scale.**

- Most of the active members of the DPM collaboration are site administrators. They do not regularly contribute to middleware development and they certainly cannot take any commitment in this sense. Partners outside CERN - e.g. national funding agencies - can try to recruit developers to contribute - if they have enough visibility on the project future - but as of today they cannot guarantee this will happen. **If DPM is to remain as one of the bases of Tier-2 storage at HL time scale, it requires an engagement from CERN to keep the leading role in DPM development**, or at least a clarification if it is not the case.
- Beside the DPM collaboration, setting up task-forces within the WLCG collaboration has proven to be highly effective to take over some activities such as integration, migration and support in a coordinated way.

DPM long-term evolution:

- Obviously, the DPM long-term future depends on decisions and actions taken now and in the next 2 or 3 years. DPM sites and the CERN DPM team have a common interest in participating as much as possible in ESCAPE, DOMA R&D, WLCG Task Forces and to share information within the DPM collaboration.
- Sites have in mind moving from RAID-6 to a JBOD setup but deployment scenarios and cost savings are not straightforward. Some options using a redundancy layer such as Ceph, HDFS and DPM on top may be considered in future.
- Making DPM suitable for future distributed storage systems can require a higher level of built-in reliability. Although the current implementation provides basic support for multiple file replicas, it still lacks important components for automatic replication and management tools necessary for simple and reliable operation. There is also some uncertainty about the technical aspects and a more sophisticated data resilience solution may be needed in the future.
- New requirements from experiments or DOMA, or further decrease of CERN support are identified as high risks for the future of DPM. In that respect, DPM sites with the help of WLCG need to be able to identify any show stoppers and to anticipate the necessity to change their plans or redefine their strategy or not. Follow-up on this should be discussed soon within WLCG.

Conclusions

A considerable amount of storage resources in a large number of WLCG sites is based on DPM as a storage management technology. In this document, the community of people interested in the DPM project described the current scenario, the future plans and the concerns for the sustainability of the DPM project in terms of manpower.

We expect to open a discussion and get some feedback from the other involved communities (CERN, EGI, WLCG, experiments and funding agencies) about their views and plans.

Acknowledgments

Special thanks go to Fabrizio Furano, Oliver Keeble, Marteen Litmaath and Matthew Nguyen for their useful discussions, remarks and corrections.

References

- [1] DPM - Disk Pool Manager : <http://lcgdm.web.cern.ch/dpm>
- [2] DPM Workshop 2019 : <https://indico.cern.ch/event/776832/>
- [3] [WLCG DOMA \(Data Organisation Management and Access\) Activities](#)
- [4] [On-going activities related to data access](#) in the context of WLCG DOMA
- [5] [DOME, a rest-inspired engine for DPM](#)
- [6] [Tokyo Tier-2 Site Report](#), Tomoe Kishimoto on behalf of the Tokyo Tier-2 operation team, HEPIX Spring 2019
- [7] [DPM @ Prague](#), Petr Vokáč, DPM Workshop 2019
- [8] [Using a dynamic data federation for running BELLE II simulation applications in a distributed cloud environment](#) M. Eber on behalf of the HEP-RC UVic group, CHEP 2018
- [9] DPM Collaboration : <https://twiki.cern.ch/twiki/bin/view/DPM/DpmCollaboration>
- [10] [WLCG DPM Upgrade Task Force](#)
- [11] H2020 ESCAPE project : <https://projectescape.eu/>
- [12] H2020 ESCAPE Kick-off meeting : <https://indico.in2p3.fr/event/18323/>
- [13] [WLCG DOMA ThirdParty Copy Activity](#)
- [14] [WLCG DOMA QOS Activity](#)
- [15] [Distributed caching system for multi-site DPM storage](#) A.Doria et al., CHEP 2018
- [16] [An http data-federation ecosystem with caching functionality using DPM and Dynafed](#) S.Pardi et al., CHEP 2018
- [17] [Integration within ATLAS of a DPM federated storage](#), S.Jézéquel et al., CHEP 2019
- [18] [The deployment of a large scale object store at the RAL Tier-1](#) A Dewhurst et al. 2017 J. Phys.: Conf. Ser.898 062051
- [19] [OSG-LHC GridFTP and GSI Migration Plans](#)

List of signers

Aurélien BAILLY-REYRE (LPNHE, Sorbonne Université / IN2P3/CNRS)
Catherine BISCARAT (CNRS/IN2P3, L2IT, formerly at LPSC)
Daniel BLOCH (IPHC, IN2P3/CNRS and University of Strasbourg)
David BOUVET (CCIN2P3, IN2P3/CNRS)
Gianpaolo CARLINO (INFN-Napoli, Italy)
Juan Carlos CARRANZA (IN2P3-CPPM / CNRS / Aix Marseille University)
Jean-Claude CHEVALEYRE (LPC Clermont, IN2P3/CNRS)
Frédérique CHOLLET (LAPP, IN2P3/CNRS)
Eric COGNERAS (LPC Clermont IN2P3/CNRS and University Clermont Auvergne)
Catalin CONDURACHE (EGI Foundation)
Alessandro DE SALVO (INFN-Roma1, Italy)

Alessandra DORIA (INFN-Napoli)
Laurent DUFLOT (IJCLab, IN2P3/CNRS and Universite Paris Saclay)
Christine GONDRAND (CNRS/IN2P3, LPSC)
Baptiste GRENIER (EGI Foundation)
Csaba HAJDU (Wigner RCP, Budapest)
Brij Kishor JASHAL (TIFR, Mumbai)
Edith KNOOPS (Aix Marseille Univ, CNRS/IN2P3, CPPM)
Tomoe KOSHIMOTO (ICEPP, U. of Tokyo)
Chun-yu LIN (TW-NCHC)
Marek MAGRYS (Cyfronet, Krakow)
Jérôme PANSANEL (IPHC, IN2P3/CNRS and University of Strasbourg)
Silvio PARDI (INFN-Napoli, Italy)
Yannick PATOIS (IPHC, IN2P3/CNRS and University of Strasbourg)
Guillaume PHILIPPON (GRIF-IJCLab, IN2P3/CNRS and Universite Paris Saclay)
Adrien RAMPARISON (IJCLab, IN2P3/CNRS and Universite Paris Saclay)
Andrea SARTIRANA (GRIF-LLR, IN2P3/CNRS, Inst. Polytechnique de Paris)
Frederic SCHAER (CEA / DRF / IRFU)
Philippe SERAPHIN (LAPP, INP3/CNRS)
Sam SKIPSEY (U. of Glasgow)
Bernardino SPISSO (INFN-Napoli, Italy)
Alessandro TARASIO, Università della Calabria & INFN-Cosenza, Italy
Enrico TASSI, Università della Calabria & INFN-Cosenza, Italy
Elisabetta VILUCCHI (INFN-LNF, Italy)
Petr VOKÁČ (Czech Technical Univ. Prague)