DPM White Paper

Please add your contribution to the different sections. Prepend it with a label in red identifying your site or institutions or region (and write in black). E.g.:

FR Sites: current deployed infrastructure…

Each section has a couple of line introducing the contents.
If you want to add something that does not seem to fit in any of the sections, add it in the "Other inputs" session.
At the end of the document you may find some useful links to retrieve some info.

Introduction / Abstract

The idea of writing a DPM (Disk Pool Manager) white paper comes out of the discussion "Looking into the future" organized during the last DPM workshop 2019 in Bern. The intent is to gather inputs from sites related to medium and long term future of DPM, share views towards the possible evolutions of our DPM infrastructures in the context of WLCG DOMA and for the HL-LHC era.

The expected goal of this White paper is to issue some forward-looking statements about DPM sustainability. Such statements, though naturally affected by uncertainties, are to reflect a common view of the DPM sites community.

?? More explicitly the question we would like to answer is: "under which conditions the community of DPM sites will consider the project sustainable and continue to invest on it." ??

How the community : DPM sites, WLCG, EGI, CERN … can make the project sustainable...

Current Status

DPM world-wide infrastructure

Over the years DPM storage system has been the preferred choice for small and medium size Tier-2 sites - mainly due to its ease of deployment and management. As to today DPM is still one of the most popular storage solutions among the community of Grid sites of both WLCG and EGI infrastructures.

More than half of those sites are part of the WLCG infrastructure, nevertheless more than 40 sites from around 15 countries (mostly in Europe but also outside) are using DPM to support the storage-related needs of various EGI-specific communities.

There are currently environ 85 DPM sites running environ 90 DPM instances - mostly WLCG Tier-2's - providing roughly 100PB of storage (for comparison, the total LHC Tier-2 pledged disk is environ 240PB). Although a big part of this community is still composed of "small" sites, there are 35 instances above 1PB providing a total of 90PB. Over the last few years the general trend has been to consolidate fewer sites with bigger - multi PB - storages. The largest single DPM instance to date is deployed by the University of Tokyo Tier-2 and it provides 10PB of disk.

DPM is widely represented in the world wide grid, indeed environ 25 countries have at least one DPM instance. National communities mostly contributing in terms of storage capacity are: UK (12 instances, 24PB), France (14 instances, 22PB), Japan (1 instance, 10PB), Taiwan (5 instances, 10PB), Italy (12 instances, 9PB). 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.

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 10TB - depending on server generations - and connection bandwidth from 1Gbps to 40Gbps per server (though 10Gbps servers are the most common case).

UK (Sam Skipsey): smaller sites moving to / encouraged by ATLAS to move to cache-only storage, using xrootd caches (simple enough, small enough to not need DPM - single server storage units with an xrootd proxy cache fronting them). Future of DPM is in the medium-large sites. All of our sites are multiple raid-6 serviers with between 1 and 10Gb/s connections. (This is not sustainable, as should be clear to anyone following technology trends.) Some large sites migrating to industry-standard storage solutions which support distributed resilience.

FR (Sartirana): 15 DPM instances in 7 sites for a total of ~22PB (over a total of 50PB of grid storage in France, 26PB at T2's). Many instances above 2PB, the largest - GRIF-IRFU - is 4.5PB. In most cases the HW deployed are simple Raid-6 servers with 10Gbps connections.

Tokyo (Tomoe Kishimoto): Tokyo is the only DPM site in Japan. DPM storage of ~10PB is reserved for the ATLAS experiment. The storage consists of 24 file servers. Each file server has two disk arrays, which are configured as RAID6.

IT (A. Doria): DPM storage system is used since 2006 in 3 out of the 4 ATLAS Italian Tier2s; the largest ATLAS Tier3 in Italy is also DPM based. All Tier2 instances are currently above 2PB. Other INFN sites, like Belle II regional center in Naples and some smaller sites, use DPM as storage system, keeping data for different projects like KM3net, PADME, and CTA. The total amount of disk space available with DPM in Italy is about 9 PB. This number is foreseen to double by the end of 2020, thanks to the ongoing national project I.BI.S.CO.. In all instances, storage HW are deployed as Raid6 arrays with 10Gbps connection to servers.

CZ (Petr Vokac): We have in production one bigger DPM storage with ~ 4PB and one small with 200TB managed by the same group. These two storages are used by ATLAS (3PB), Auger (0.4PB), DUNE (0.3PB) and Belle (0.1PB) experiments. Whole storage space is provided by 15 servers with appropriate number of disks in each RAID6 group to ensure reasonable reliability. Older machines

have 10Gb connectivity, but most recent purchase from 2018 use 40Gb networks. Our site usually add ½ of the existing storage capacity every ~ 3 years and new purchases will be done next year.

EGI (Baptiste Grenier): DPM is used by many sites part of the EGI infrastructure, more than half of those sites are part of the WLCG infrastructure, nevertheless more than 40 sites from around 15 countries (mostly in Europe but also outside) are using DPM to support the storage-related needs of various EGI-specific communities. EGI has been testing and distributing LCG-DM/DPM as part of his middleware distribution for more than 10 years and DPM is currently the most reasonable solution for sites looking for an easy to use and stable while still performant and scalable storage solution. EGI is providing some infrastructure services useful to the DPM community and users: UMD, GOCDB, ARGO, APEL and GGUS are some examples of those services. DPM is available via UMD repositories, and staged rollout is in place for it, to be able to validate DPM releases in the EGI environment. DPM endpoints are registered in the GOCDB and so that they can be discovered and monitored via things like ARGO or WLCG-specific IS or monitoring. APEL is used to collect accounting information and usage information can be collected using the EGI Accounting Portal. GGUS is also supporting a dedicated Support Unit for DPM Development. We also have been doing coordination of updates across the federation and also have been supporting and interacting with user communities.

DPM as storage technology, current status

What is your feedback about the current status of the DPM software ; see DPM status, Fabrizio Furano@DPM Workshop 2019 (link below) ? We focus here on DOME ignoring the legacy stack and migration issues. Do you see limitations for what concerns operations today and in the middle term (5 years horizon)? Plans (and motivations) to migrate or evaluate a different storage technology?

All DPM sites (to be discussed) (Frédérique Chollet 1rst attempt to introduce the topic and summarize the contributions) :
Are DPM sites running the new DOME stack confident for the next 5 years ?

The reengineering done in the past 2 or 3 years by the DPM development team at CERN has been successful in providing the new DOME software stack and overcoming the performance limits of the SRM implementation. The gain in stability and performance has been confirmed by Tokyo and Italian sites. So it is agreed among sites that dropping SRM protocol and moving to the DOME configuration is the way to go.

Intensive stress tests performed at Prague Tier 2 have shown that the DOME configuration increases dramatically the performances. The gain is different for gsiftp, xrootd and webdav protocol, but metadata operations are at least an order of magnitude faster compared to SRM. In Prague, the current storage hardware utilization is on average below 10% with respect to the numbers reached during the DOME stress tests and it has been demonstrated that even with such light diskservers utilization the legacy stack running SRM can reach its limits serving ~5k cores. On the other hand, this makes us confident that from a performance point of view, DOME can deal with transfers for a typical 8k core T2 site in the future (even with the same headnode / database performance and our usual 15% cpu capacity increase a site should be safe for at least a decade considering similar workloads).

A.Doria SRM stack and DOME stack are completely independent. DPM sites proved to be able to easily and effectively switch to SRM-less mode. Are the other current storage technologies ready to switch off SRM in such an effective way?

From Tokyo experience, we know that DPM is able to scale up to 10 PB or more.

So DPM should remain an appropriate storage technology for Tier2 sites unless requirements from experiments or DOMA major changes .

DPM is fine with TPC and seems to be fine in developing new token based access. DPM is part of DOMA TPC and QOS activity.

IT (A. Doria): After the migration of the DPM software to the new DOME flavour, the system has gained higher stability and performance, fulfilling all the requirements of the sites and of the involved experiments at the moment. We consider current DPM release capable to support storage evolution in Italian Tier2 sites in the short/medium term. For the long term (HL-LHC), we could consider moving to other technologies only in case any critical WLCG-DOMA requirement could not be fulfilled/implemented in DPM, as this change would imply a considerable effort in terms of manpower at the sites.

Tokyo (Tomoe Kishimoto): Tokyo had a performance/scalability issue of SRM, but it has been solved after migrating to the DOME and changing the main transfer protocol in ATLAS. Minor issue: Tokyo see more than > 1M error messages every day. I don't think all of them are serious errors. An improvement of log severity/message will be helpful for operation.

CZ: Our storage server hardware utilization is on average below 10% with respect to the numbers reached during stress tests. With DPM SRM implementation our storage reached its performance limits even with such light diskservers utilization. Dropping SRM protocol for majority of our transfers and moving to the DOME configuration increased dramatically performance. Our DOME stress test showed that performance gain is different for gsiftp, xrootd and webdav protocol, but metadata operations are at least an order of magnitude faster compared to SRM. This makes us confident that from a performance point of view DPM can deal with transfers for our 8k core T2 site in the future (even with the same headnode / database performance and our usual 15% cpu capacity increase we should be safe for at least a decade considering similar workloads).

EGI (Catalin Condurache): EGI has recently started a campaign to migrate and enable DOME across the 40 sites providing DPM storage to various EGI communities. It is proceeding without any particular issues, with most sites finalising the upgrade, few of them being delayed by local related issues.

DPM collaboration, current status

This section is dedicated to the DPM collaboration. How it was initially setup ? How it works ? Which are the contributions to DPM development and support ? make a reality check based on available resources... what is it capable to sustain ? is it enough ?
Let us try to identify a sustainability mode for DPM and the investment needs in development and support for the next 5 years.

The DPM collaboration 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 features implementation. It is perhaps useful to recall that the development of the new DOME stack 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 necessary effort for possible large scale new development projects.

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. But it is 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 available/first released in EPEL-testing repositories. It is important to quote 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 in the EPEL-testing activity on testbeds and advanced features validation (multisite 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) has contributed a lot in DOME validation including debugging and problem solving.
- *Any missing contributions should be added here...*

It is worth noting that EGI has not been expressly part of the DPM collaboration but has been testing and distributing LCG-DM/DPM as part of his middleware distribution for more than 10 years. DPM is considered as the most reasonable solution for sites looking for an easy to use and stable while still performant and scalable storage solution. EGI is providing some infrastructure services useful to the DPM community and users: UMD, GOCDB, ARGO, APEL and GGUS are some examples of those services. DPM is available via UMD repositories, and staged rollout is in place for it, to be able to validate DPM releases in the EGI environment. DPM endpoints are registered in the GOCDB and so that they can be discovered and monitored via things like ARGO or WLCG-specific IS or monitoring. APEL is used to collect accounting information and usage information can be collected using the EGI Accounting Portal. GGUS is also supporting a dedicated Support Unit for DPM Development. We also have been doing coordination of updates across the federation and also have been 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 and made it 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 have played in favor of DPM being a healthy community.

Nethertheless, there is no certainty that the DPM collaboration 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.

Beside the DPM collaboration, the WLCG DPM Upgrade Task Force 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 ad-hoc task force within WLCG and this working mode may help in future as well.

For the next five years, some of the upcoming technical evolutions are known and some operational concerns are expressed by sites...
Is the DPM collaboration as it is, capable to cope with that ? Statements about possible evolutions regarding the collaboration :
Beyond that, is the DPM collaboration as it is, capable to address the long term challenges ?

Ongoing R&D activities

Which are the R&D activities ongoing in your site/region concerning DPM storage? And those foreseen in the middle term scale (5 years)? Which are the motivations?

UK (Sam Skipsey): centralised banning/ARGUS integration.

FR (Frédérique Chollet): Four French sites have put their expertise in common and set up a testbed providing a single distributed DPM instance, fully integrated in the ATLAS Grid infrastructure. The integrated DPM caching mechanism based on DPM volatile pool is implemented (currently looking for ATLAS use cases of an integrated cache, read-ahead functionality expected soon or later). This infrastructure is foreseen to be a testbed for a

DPM component within a Datalake (ESCAPE/DOMA prototype). This is rather in line with IT approach.
IT (A. Doria): The keywords of activities about DPM in Italy are “distributed storage” and “caching”. Using DPM technology we deployed a test-bed configuration where storage pools are distributed among three Italian sites and accessible through a single endpoint. The volatile pools provided by DPM sw were used to implement a cache area in each site. The encouraging results of this activity have been reported at CHEP 2018 (https://www.epj-conferences.org/articles/epjconf/abs/2019/19/epjconf_chep2018_04056/epjconf_chep2018_04056.html).

This established the basis to continue R&D about the use of caching with DPM in the ESCAPE project, in which INFN-NAPOLI and INFN-ROMA1 are involved. (new) Furthermore, INFN sites in ESCAPE intend to test, in the DataLake context, DPM built on a CEPH file system and distributed DPM with http TPC..

CZ: DOME DPM supports only GLUE2 BDII schema and it seems EGI ARGO monitoring is currently not completely happy with missing GLUE1.3 (is this schema officially deprecated?). Centralized banning depends on legacy DPM code and the design is a bit problematic from a security point of view, because only users that already accessed DPM storage in the past are considered for blacklisting. Integration of the lcmaps with DPM doesn’t seems to be easy, because of missing user data (certificate issuer). Also lcmaps/ARGUS brings additional dependency on packages that are not available in EPEL. I’m not sure how much effort should be spent on this issue, because with token based authentication this mechanism will become obsolete.
According “GSI Migration Plans” support for GSI can start to disappear pretty soon (2021) and unless there is token -> x509 translation service available storage implementations should support for WLCG/SciTokens. Token based access can become critical requirement for storage that would like to participate in the distributed data store.
Running consistency checks (dpm-dbck) should be possible without downtime.

Operations, RH, investments

which are your investment plans (up to HL-LHC scale) for the storage (manpower, HW, ...) and how they impact or are impacted by the choice of storage technology and the evolution of DPM?

FR (Frédérique Chollet) : The increase of storage usage at HL-LHC horizon will also induce scalability challenges on the storage operation by site administrators. So French sites are considering the consolidation of DPM end-points, mutualisation of critical services and expecting a reduction of operational complexity and cost.

FR (Frédérique Chollet) : Two concerns related to investments : Hardware purchases are now made with 7 years warranty (!). A change of technology requires peak of investment whereas the French sites work with annual flat budget.
IT (A. Doria): , In terms of HW , the sites that currently deploy DPM storage in Italy are not supposed to grow in a way that cannot be managed with the current technology. HW acquisition are normally made on a flat budget basis. But next year the I.BI.S.CO national

project will provide about 10PB of raw disk space to Naples site in Raid6 technology. At the same time, manpower will hardly increase, so we our needs go in the direction of using consolidate solutions and reducing operational costs. Coordination among the sites, use of shared services and possibly building a distributed system among sites could help.

CZ: We usually buy storage HW every ~ 3 years and new purchase usually comes with significant increase of the storage capacity. Assuming similar situation in the future we can probably delay our decision till 2023/24. If DPM doesn't fit HL-LHC requirements it'll be much more clear in 2024 which storage technology is mature enough with low manpower requirements for deployment and operation (dropping SRM and GridFTP could make different storages more attractive).

Views towards the future (and site concerns)

What is your long time view for your storage. ? Which are you doubts/questions?
Which is the sustainability of the DPM sites, DPM collaboration/project to HL-LHC scale? is there a strong community in place? a prospective for a strong collaborations? which may be the expected contributors?

Sites provide a service to LHC computing. So the support CERN gives to DPM translates in PBs of operated storage for WLCG. Discontinuing this support may have an impact on this service.
FC How far can DPM scale-out by adding data pools spread across sites (within a few ms in RTT) to a single headnode ?
FR (Sartiana) Commenting/expanding on FC input (not clear if this goes in this section or the next). The possibility to scale to big DPM pools, possibly geographically delocalized, is an interesting subjects. IMHO one key point is the possibility to have "horizontal" scaling of the head node. The main constraint in big pools is that the DB becomes huge and quickly unmanageable. One possibility is to have a "multi-head" DPM installation with different databases taking care of different subtrees of the namespace (e.g. subtrees of the different VO's basically never interact...). Alternatively the whole DPM DB has to be re-thought to make it way more scalable. HA on the head node HA is another important point if we want to have big distributed pools.
A.Doria HA headnode could be added in the needed features. I agree that a huge DB is one of weakness of a very large and distributed system.

Tokyo (Tomoe Kishimoto): Scalability/performance of the DPM is the biggest concern of Tokyo site. As Sartiana commented, high availability is also an important point for us to avoid downtimes due to DPM or server maintenance..

CZ: We currently use DPM, pure XRootD (ALICE) and small dCache + in the future there may be new requirements to provide object storage (e.g. DUNE is considering object store advantages/disadvantages for their computing model). We have to make that list shorter…
I'm not so concerned about headnode HA, because DB can be configured with replication and headnode is just simple virtual machine that doesn't basically keep any state - I think simple DPM headnode "active/passive" cluster could be configured even now although I never tried to test such configuration. Support for splitting database for /dpm/fqdn/home/vo1 and /dpm/fqdn/home/vo2 would make DPM DOME even more scalable (same DPM pool sharing transfers for all VOs distributing load on all diskservers, but with per-VO DB), but unless experiment workflow significantly changes - e.g. use (ask for metadata) order of magnitude more files for a job - my stress tests makes me quite

confident that one DB server with reasonable hw can deal with load from DOME DPM. Still I have to admit this looks like an interesting idea.
Still it would be nice to know which way of HA configuration makes sense from developers point of view.

EGI (Catalin Condurache): EGI is consistently involved in various use cases for EOSC-hub related projects (Competence Centers) and these are usually requiring access to compute and data storage provided by participating sites.
Federated data has become more and more a 'must' requirement and suitable technologies and solutions are being explored. We consider DPM a good candidate for storage technology within Data Federation, also offering the users a standard interface (WebDav) for access and Third Party Copy (TPC) as a more efficient transfer mechanism, and in the years to come, as part of participation to various European funded projects, EGI is willing to setup and run test pilots that would make use of DPM.

Needs, Innovative approaches, Alternatives studies, Impact of a change of technology

Proposals for innovations and possible evolutions and how they are motivated by your operational needs as a site.
One key point is that we do not yet know what DOMA will ask from us in the end. We may write a list of things which are "in the air". But it is not worth to enter in the details.

IT (A. Doria): *Improving DPM caching mechanisms seems one of the interesting points~~. Our participation to ESCAPE project is especially oriented to explore how DPM caches could be integrated in a Data Lake. We are trying to find in ESCAPE some funding to have a person who could work in this field, with a direct connection to the DPM collaboration~~. (changed, see "ongoing R&D activities")*

Interest in Object Storage is growing, a support in DPM for this kind of underlying technology would be welcome.

FR (Sartirana) one point on DPM over different backend (Sam evoked the DPM over HDFS in a commentary below). This may be an interesting subject about future developments but I'm just a bit confused on this (developers may comment). Currently DPM is an implemented posix-like distributed file system backend plus a set of standard communication protocols (xrootd, gsiftp, https). The key feature is that the implemented backend is very a simple file-based distributed filesystem, with somewhat limited fonctionnalites but very easy to manage. Which are the scenarios for different backend solutions? Which are the advantages?

CZ (P.V.) DOME design supports plugin architecture for backend storage, but only posix filesystem based plugin is currently implemented. Only one site was using DPM with HDFS and that did not justify the effort of DPM developers to support this backend in the DOME.

EGI (Catalin Condurache): Non X.509 (token) based access is another 'must' 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).

Current views - Statements of the future

Potential requirements - List and cost (in development) of additional functionalities

Requirements that may become critical

- **Central banning / ARGUS integration**
- **Support for WLCG/SciTokens** Token based access can become a critical requirement for storage that would like to participate in the distributed data store.

Less critical (potential requirements or wish list)

- **Support for GLUE 1.3 (EGI ARGO)**
- **Storage consolidation** (increase of the storage capacity / headnode >> 10 PB ??) **:**
    - **DPM scale-out** by adding data pools spread across sites (within a few ms in RTT) to a single headnode
    - **High availability**

List of functionalities that are in line with emerging technologies and advanced features (a.k.a datalake ones)

- **Improving DPM caching mechanisms :** The participation of Italian DPM sites to ESCAPE project is especially oriented to explore how DPM caches could be integrated in a Data Lake
- **DPM over different backend solutions with new resilience profiles :**
    - Support for object storage
    - JBOD setup instead of RAID-6
- **QOS**

Statement on the sustainability of DPM

Other inputs

Put here everything which does not fit in the previous sections (including suggestions for adding new sections). This is also a good place for EGI and DPM dev to add their contributions

Prague :

1. outputs of DOMA activities are unknown for now - some requirements can be critical for storage to be able to participate (e.g. WLCG/SciToken support) other may be less critical for smaller sites (QoS?)
2. No motivation to migrate to the different storage technology at least till HL-LHC and then reevaluate how DPM features fits DOMA requirements for site of our size
3. performance / scalability limits - legacy DPM / SRM had hard time to serve our ~ 5k cores, DPM DOME performance at least an order of magnitude better - fine for our site for a long time

Laurent Duflot : Should DPM support QoS?
Also most sites have Raid-6 storage and one way to increase capacity is to pass to a JBOD setup. This requires resilience to disk loss and effective tools to identify the portion of data lost and to notify it to central operations. Which would be the adaptability of DPM to this type of setup.

From DOMA-Access: also the cache setup is a possible solution, which are the perspectives of cache functionalities in DPM ?

Various links :

1. Minutes of DPM white paper kickstart meeting : https://indico.cern.ch/event/847892/
2. WLCG Site Survey : https://twiki.cern.ch/twiki/bin/view/LCG/QoSSurveyAnswers
3. (storage part) and conclusions https://twiki.cern.ch/twiki/bin/view/LCG/QoSSurveyConclusions
4. DPM status, Fabrizio Furano@DPM Workshop 2019 : https://indico.cern.ch/event/776832/contributions/3378511/attachments/1862155/3060730/StatusDPM2019.pdf
5. WLCG strategy document : http://wlcg-docs.web.cern.ch/wlcg-docs/technical_documents/WLCG%20Strategy%20towards%20HL-LHC.pdf
6. Data-lake strawmann model https://docs.google.com/document/d/1WslZOavtI1wruWzzpwIyqm3AcIeaGRiH1FbAFWFceeo/edit#heading=h.wsyv7oxoht0k
7. List of on-going activities related to data access in the context of WLCG DOMA. https://bit.ly/2wGQ2K6
8. https://twiki.cern.ch/twiki/bin/view/LCG/DPMupgrade#GOCDB_DPM_reachable_by_srmPing_o
9. DOME A rest-Inspired engine for DPM : https://gitlab.cern.ch/lcgdm/dmlite/blob/master/doc/dome/dome.pdf
10. DOME metadata performance - https://indico.cern.ch/event/776832/contributions/3378514/attachments/1861982/3060425/dpm_prague_vokac.pdf - slide 6 - table with number of metadata operations (read/write/stat/delete) per second for all protocols (SRM/GridFtp/XRootD/WebDAV) for different number of concurrent clients (threads)

11. List of DPM sites sorted by size https://twiki.cern.ch/twiki/bin/view/LCG/DPMupgrade#GOCDB_DPM_reacha
12. DPM UK status 2016@DPM Workshop 2016 (Whislist and Future Look T2s https://indico.cern.ch/event/559673/contributions/2268618/attachments/1375996/2089440/DPMWorkshop2016UK.pdf Sam Skipsey on behalf of GridPP Storage Group
13. WLCG Authorisation Working Group: https://twiki.cern.ch/twiki/bin/view/LCG/WLCGAuthorizationWG