



LUND
UNIVERSITY

About LDCS

Lightweight Distributed Computing System

FLORIDO PAGANELLI, RESEARCH ENGINEER, LUND UNIVERSITY
FLORIDO.PAGANELLI@FYSIK.LU.SE




Outline / Purpose of this talk

- What is LDCS
- Challenges in building distributed computing system for small communities
- Possible synergies within the community

Purposes of the **L**ightweight **D**istributed **C**omputing **S**ystem




Run LDMX software [3]
computing jobs in
containers on a **distributed**
infrastructure



Orchestrate
submission
campaigns of
200K-300K jobs

Manage **distributed data storage**



Expected
total **10PB**

Thousands of
files **~100MB**




Respect **authorization** and
authentication rules across
borders



Motivation for LDCS development

- LDMX is a “**small**” **community** (~40 people) but...
 - The computing needs of LDMX are **comparable to the ones of other Particle Physics projects based at CERN** that run on the Worldwide LHC Computing GRID (WLCG) [3]
 - Downsides:
 - Small research communities such as LDMX **cannot put in the same time and effort** that CERN invested in WLCG
 - Unless running a CERN related project, it becomes quite unrealistic for a small project to fulfil all the requirements needed to utilize the WLCG
 - Upsides:
 - WLCG-alike expertise is available among project members or institutes
 - the products involved are open source and can scale if properly configured
- => **Challenge:** can we build a system that is *like* WLCG, but less complicated -- **lightweight?**







Strategy and approach



- **Reuse in-home experience** with WLCG GRID [5] or the like.
 - Lund is part of NorduGRID [4] that developed the ARC Computing Element, a service that manages jobs in WLCG clusters
 - Other institutes are also running WLCG components for existing CERN projects, that can be reconfigured for smaller needs.
- **Reuse existing computing resources** that users from the participating users have access to – National HPC clusters or local university clusters – even if these do not have the capabilities to be part of WLCG
 - Adapt products to solutions already existing in the participating institutions
 - **Overcome border authorization limitations, while still following the rules**
 - Reuse knowledge and personnel at the sites so that the **effort** required by local administrators and researcher is **minimal**


Products involved in the implementation



see 2021 whitepaper "Building a Distributed Computing System for LDMX" [1] and LDCS software repository [2]
Credit to L.K. Bryngemark, B. Konya, D. Cameron, G. Mullier, A. Salnikov, , et al

How does it work

7/29




LDCS workflow

- 1) The researcher sends a computational job request to **aCT** using the **aCT client**
- 2) **aCT** does *data matchmaking* and *job submission* by asking **RUCIO** where the data is geographically located, and then submits the job to the corresponding cluster
- 3) **RUCIO** maintains a *catalog* with information about where the data is located
- 4) **ARC** provides access to different clusters for users belonging to different organizations
- 5) **Grafana monitoring** and **aCT job logs** provide a status dashboard
- 6) **Researchers** retrieve their job results via **RUCIO** or directly from the clusters by using **ARC tools** or via the cluster's own storage interface

Feat_hures

8/29


Job Description




Automatic data movement
Automatic cleanup




Data model




Easy setup
For datacenters




Data Replication



Automatic resubmission
cluster status check




Cluster Throttling



Limitations

9/29

- **Single submission point.** At present there is no *remote client* developed. The submission must happen via a trusted machine where aCT is installed.
- **Authentication/Authorization:**
 - On the infrastructure/administrator side, **certificate distribution must be done manually once a year.**
 - On the user side, **only those LDMX physicists who have access to the aCT machine** are allowed to submit and query RUCIO database to find where data is.
- **Key expertise is required in case of system failures.**
- Certain datacenters do not accept “edge services” such as ARC due to **internal security policies... and politics**



Operations 1/4

10/29

We ran a pilot to measure workload in 4 data centers, 5 institutions contributing to different tasks/roles with **existing** personnel, on tasks they **were already doing anyway**, sometimes within WLCG.

- Lund university, Sweden
 - **Technical researcher** and **LDMX Physics researcher** runs coordination of meetings, if needed and when needed
 - **LDMX Physics researcher** plans job campaigns, constructs job specification
 - **System administrator** maintains ARC and Gridftp services
- Oslo University, Norway:
 - **aCT developer** maintains orchestration and data management infrastructure (aCT, Rucio)
 - **LDMX Physics researcher** from Lund has access to the submission infrastructure
- SLAC, USA
 - **Local cluster administrators** manage ARC, allocate existing GridFTP storage for LDMX
- Caltech, USA
 - **Local cluster administrators** manage ARC
- UCSB, USA
 - **Local cluster administrators** manage ARC

Operations 2/4 – workloads

11/29

- Initial installation and deployment was coordinated by Lund. The configuration files and instructions were provided to the system administrators.
- ✓ – Administrator reported that installation was smooth and did not require long time.
- Tests were run by the **LDMX Physics expert** and the **aCT developer in Oslo independently**.
- ✓ – This **did not require any action from the system administrators**, unless some problem was discovered.
 - We organized bi-weekly meetings of **~1h-1h:30m** to check that all parties were happy with the current status or brainstorm or help solving problems that arose during the campaigns.
- While in operation the system components do not need any human intervention except:
 - ✓ – Preparation for certain campaigns required intervention from system administrators, who reported a maximum of **1h** work
 - Ordinary maintenance such as update of the product or certificates (once or twice a year)
 - Lack of disk space on the target system, or other local situations that would prevent jobs to run as planned – **up to three times a year to the record**

Operations 3/4 2021-2024

12/29

bi-weekly meetings to solve technical issues or discuss downtimes
Cancelled if no problem – we cancelled many, **low overhead**



- 1 large campaign per year,
~ **200 000** Jobs lasting 3-4weeks
- Several smaller campaigns for other tasks



To compare: the COSMOS cluster in Lund allows for maximum submission of **300 jobs** at time and each job can run maximum **one week**

Total storage used: ~**300TB**

(~**600TB** counting replication on both storage servers)

Average file size of single file produced by jobs: ~**100MB**

Calculations contributed to the following scientific output:



- *Photon-rejection power of the Light Dark Matter eXperiment in an 8 GeV beam*: [https://doi.org/10.1007/JHEP12\(2023\)092](https://doi.org/10.1007/JHEP12(2023)092)
- *LDMX Design Report*: <https://arxiv.org/abs/2508.11833>



Grafana dashboard

13/29

EDMA jobs



number of jobs in different states over time during a simulation campaign in the second half of December 2020. The different colors represent different batches of 5000 jobs which were submitted over time.

Operations 4/4

14/29

- **Paused in 2024:**

- less need for massive computational needs
- aCT developer changes within NorduGRID
- one computing center could not commit anymore



- **Getting ready for new campaigns 2026/2027**

- All system restored to latest versions of different components
- Plan to experiment with other storage protocols supported by ARC



What is "lightweight" exactly?!?

15/29

Lightweight **hardware requirements**

Software runs on any hardware
(usually VM)

Researcher does **not need to refactor**
Computational job matches LDMX workflow

Quick installation and low maintenance for system administrators






No need to comply with many **lengthy policies** as in WLCG

No ticket system: direct contact with developers/architects

Lessons (re?)learned 1/4

16/29


What are we in control of?



Lessons (re?)learned 2/4

17/29





What are we in control of?



Lessons (re?)learned 3/4

18/29




What are we in control of?



Lessons (re?)learned 4/4

19/29

What are we in control of?



Lessons (re?)learned 4/4 (text)

20/29

- Computing models developed by a certain community **tend to resist** – can't really throw away many years of coding
- Components with stable development and good amount of maintainers are safe. But:
 - External components that may be **driven by a different user base** or with different economical objectives. One needs to follow how their development progresses.
 - Rucio follows the needs of CERN experiments, hard to keep up with their changes
 - Grafana is not really targeting academia
 - We ran our own storage protocol based on Globus libraries. But they have been retired, data centers might not accept them anymore for security reasons
 - Globus has decided to change their business model, costs might be unfeasible
 - Other solutions to test such as WebDAV
- Components that depend on **only one developer** are always at risk.
 - If the developer leaves it is hard to go forward.
 - Good as long as one can develop against open source code
- Client components with small user base are very hard to maintain.
 - UI is hard!
 - Accommodating for many use cases quickly increases the complexity – maybe AI a solution?
 - Clients need to authN/authZ, see below
- Authentication and authorization is often **not something we can directly control**
 - Depends on national and international regulations e.g. NIS2 cybersecurity directives incoming
 - Requires establishing **trust** among parties – we generate our own certificates, and all involved partners accept them
 - Modern technologies require a whole substack of services that research group cannot maintain, and then they depend on external services for operations
 - Switching to **token-based** authentication requires **more complicated** rules than X509 certificate issuers from the GRID times, and very complicated intra-service communication tools and rules
 - Affects negatively also the development of clients

Open questions

- 1) Ask me more about LDCS™!
- 2) What is your experiment using at the moment?
- 3) Is your experiment facing the **same problems**?
- 4) Are there problems we could **solve together**?
- 5) Are there **components or common knowledge** that we are developing that could be useful to all of us?
- 6) Can we achieve similar levels of interoperability of CERN projects, preserving the needs and identities of each project?
- 7) Can we create a **lighter administrative/management model than CERN/WLCG**?
- 8) Are there opportunities for a **common infrastructure toolkit for DM computing**?
- 9) **Other examples of successful communities?**

References

- [1] LDCS whitepaper <https://doi.org/10.1051/epjconf/202125102038>
- [2] LDCS repository <https://github.com/>
- [3] LDMX-Software/LDCSLDMX computing software repository
<https://github.com/LDMX-Software/lmx-sw>
- [4] NorduGRID consortia <https://www.nordugrid.org>
- [5] World LHC Computing Grid <https://wlcg.web.cern.ch/>
- [6] Rucio data management framework <https://rucio.cern.ch/>
- [7] NorduGRID ARC Computing Element <https://www.nordugrid.org/arc/arc7/>
- [8] NorduGRID aCT workload manager <https://github.com/ARCCoNTrolTower/aCT>
- [9] Grafana <https://grafana.com/>

Images from OpenClipart gallery and respective software providers or institutions



LUND
UNIVERSITY

Backup slides



Purpose of the **Lightweight Distributed Computing System**

- Run LDMX software [3] computing jobs in **containers** on a **distributed** infrastructure
- Allow for **massive simulation campaigns of at least 200K jobs**
 - Provide **job orchestration tools**
- Manage **massive distributed data storage of the order of 10Pb**
 - Store data on **multiple sites**
 - Store **metadata** to identify time of computation, details about simulation/analysis output, location of stored data
- Provide a unified solution to limited **access to computing resources** (e.g. HPC clusters) between EU and US
 - Based on certain accepted/feasible access rules across international borders

Products involved in the implementation



- **C++, Python** Computing jobs running in **containers**
 - Identify computing centers that allow for **docker** or **apptainer**.



- Allow for **massive simulation campaigns**



- **Orchestration tools** required

- Use **NorduGRID ARC computing element** [7] as edge service on computing centers
- Use **NorduGRID aCT** [8] as workload and submission manager

- Manage **massive distributed data storage**

- Store data in multiple sites:
 - Identify protocols already in use: **gridftp** (e.g. Globus)
- Store **metadata**
 - Use CERN's **Rucio** [6] as main catalog manager



- Provide a unified solution to **access to computing resources** (e.g. HPC clusters) between EU and US
 - Identify an authentication and authorization framework that is trusted by centers: **old school x509 certificates**




- **Monitoring** provided by **Prometheus + Grafana** [9]

see 2021 whitepaper "Building a Distributed Computing System for LDMX" [1] and LDCS software repository [2]

What is "lightweight" exactly?!?

- All the involved software components **run on any hardware with a linux system**. No special hardware was procured for the implementation. **We reused existing hardware already operating in the facilities**. We reused the facilities means of managing data.
- **Simple computational model**. aCT job description respects LDMX own computing model
- **Reduced bureaucracy**: being formally part of WLCG has a formal path which requires time and involves project representers. In this case we just used the technical connections within the LDMX community.
- **Reduced burden on system administrators**: Installation went smoothly. I personally helped numerous WLCG system administrators setting up ARC for WLCG according to the different sources (CERN project, rules from governmental bodies, EU projects requirement. There are more rules and subsystems to configure in a real WLCG environment – all this is gone within LDCS.
- **Centralized technical coordination between Lund and Oslo**: since the software providers reside in these two universities, we were able to quickly resolve issues. No need for tickets, waiting times or complicated change request procedures.



- **Job description** indicates what data is required and where the output should be stored.
- A solid **data model** allows for precise metadata to be attached to output files.
- Data **replication** can be achieved using RUCIO.
- aCT does **automatic resubmission** and **cluster status check**. If a cluster is not operational, another one will be checked. A job is tried a certain amount of time defined by aCT configuration.
- **Results are automatically sent to the target storage** by ARC, as selected by the submitting user. If a local copy is required, the job data is retained; otherwise aCT takes care of **cleaning up the remote workspace** to save storage.
- **Cluster throttling: optimized use of the computing resource**. If the computing facility has a cap limit in the job queue, ARC can feed the queue accordingly. This allows for a queue longer than the datacenter queue, still respecting the local datacenter rules of conduct.
- Datacenter setup has proven to be an easy task to do for system administrators.

- **Single submission point.** At present there is no *remote client* developed. The submission must happen via a trusted machine where aCT is installed.
 - aCT is a complex tool, not suitable to be managed by end users. Running it requires knowledge of the software. We provide a basic client interface for submission and monitoring.
- **Authentication/Authorization:**
 - **On the infrastructure/administrator side, certificate distribution must be done manually once a year.** Once a year, Lund takes care of generating and distributing certificates as well as configuration of all the components in case of important changes.
Recent development in ARC allow to use Letsencrypt and token-based access, which we are considering.
 - On the user side, **only LDMX physicists who have access to the aCT machine** are allowed to submit and query RUCIO database to find where data is.
 - Currently, data retrieval credentials depend on the authentication schemes supported by the protocols used by each storage at every data center.
- **Key expertise is required.** If something goes very bad, usually it's the technical experts that need to fix.
- **Certain datacenters do not accept services such as ARC due to policies.** For these NorduGRID is trying to identify a solution that follows those policies.