

Table of Contents

WLCG Operations Coordination Minutes - 29 August 2013.....	1
Agenda.....	1
Attendance.....	1
News.....	1
Middleware news and baseline versions.....	1
Tier-1 Grid services.....	1
Storage deployment.....	1
FTS deployment.....	3
LFC deployment.....	3
Other site news.....	3
Data management provider news.....	3
Experiment operations review and plans.....	3
ALICE.....	3
ATLAS.....	4
CMS.....	4
LHCb.....	5
Task Force reports.....	5
CVMFS.....	5
SHA-2 migration.....	6
SL6.....	6
Tracking Tools.....	7
FTS-3.....	8
XrootD Deployment.....	9
gLExec.....	9
AOB.....	9
Action list.....	9

WLCG Operations Coordination Minutes - 29 August 2013

Agenda

<https://indico.cern.ch/conferenceDisplay.py?confId=263201>

Attendance

- Local: Maria Girone (chair), Andrea Sciabà (secretary), Simone Campana, Joel Closier, Stefan Roiser, Felix Lee, Steve Traylen, Oliver Keeble, Nicolò Magini, Ian Fisk, Ikuo Ueda, Maarten Litmaath, Andrea Valassi, Markus Schulz, Alberto Aimar, Michail Salichos, Alessandro Di Girolamo, Domenico Giordano, Helge Meinhard, Maria Dimou
- Remote: Sang-Un Ahn, Tommaso Boccali, Andreas Petzold, Michel Jouvin, Alessandro Cavalli, Jeremy Coles, Di Qing, Massimo Sgaravatto, Oliver Gutsche, Rob Quick, Peter Solagna, Isidro Gonzalez Caballero, John Kelly, Alessandra Forti, Salvatore Tupputi, Peter Gronbech

News

Maria announces that the job/machine features task force is now fully set up.

The also announces that from January 1st she will start working as computing coordinator for CMS and she will begin on October 1st to overlap with Ian Fisk. The next WLCG operations coordination meeting will be the last she will attend as chair.

Middleware news and baseline versions

<https://twiki.cern.ch/twiki/bin/view/LCG/WLCGBaselineVersions>

Highlights:

- new BDII release in the latest EMI-2/3 update, including better GLUE-2 support and security fixes. Sites should update all their BDII instances
- new CVMFS version released for a security fix. Sites should upgrade or at least apply the hot fix in the above twiki
- perfSONAR: sites should upgrade to the latest version, fixing many deployment problems

The end of support for dCache 1.9.12 has been postponed to September 30 due to a delay in releasing the SHA-2 compliant version in the dCache 2.2 series.

Tier-1 Grid services

Storage deployment

Site	Status	Recent changes	Planned changes
CERN	CASTOR: v2.1.13-9-2 and SRM-2.11 for all instances (SRM-2.11-2 for LHCb) EOS: ALICE (EOS 0.2.37 / xrootd		

Site	Status	Recent changes	Planned changes
	3.2.8) ATLAS (EOS 0.2.38 / xrootd 3.2.8 / BeStMan2-2.2.2) CMS (EOS 0.2.29 / xrootd 3.2.7 / BeStMan2-2.2.2) LHCb (EOS 0.2.29 / xrootd 3.2.7 / BeStMan2-2.2.2)		
ASGC	CASTOR 2.1.13-9 CASTOR SRM 2.11-2 DPM 1.8.6-1 xrootd 3.2.7-1	None	None
BNL	dCache 2.2.10 (Chimera, Postgres 9 w/ hot backup) http (aria2c) and xrootd/Scalla on each pool	None	None
CNAF	StoRM 1.11.2 emi3 (Atlas, LHCb) StoRM 1.8.1 (CMS)	Atlas-LHCb update	CMS update: september
FNAL	dCache 1.9.5-23 (PNFS, postgres 8 with backup, distributed SRM) httpd=2.2.3 Scalla xrootd 2.9.7/3.2.7.slc Oracle Lustre 1.8.6 EOS 0.3.1-5/xrootd 3.3.3-1.slc5 with Bestman 2.2.2.0.10		
IN2P3	dCache 2.2.12-1 (Chimera) on SL6 core servers and 2.2.13-1 on pool nodes Postgres 9.1 xrootd 3.0.4		
KISTI	xrootd v3.2.6 on SL5 for disk pools xrootd 20100510-1509_dbg on SL6 for tape pool dpm 1.8.6		xrootd upgrade foreseen for tape (20100510-1509_dbg -> v3.1.1) in September
KIT	dCache <ul style="list-style-type: none"> • atlassrm-fzk.gridka.de: 2.6.5-1 • cmssrm-fzk.gridka.de: 2.6.5-1 • lhcb-srm-kit.gridka.de: 2.6.5-1 xrootd <ul style="list-style-type: none"> • alice-tape-se.gridka.de 20100510-1509_dbg • alice-disk-se.gridka.de 3.2.6 • ATLAS FAX xrootd proxy 3.3.1-1 	dCache upgrade to 2.6.5 end of July	alice-tape-se.gridka.de will be upgraded to xrootd 3.2.6 in Sep/Oct. The ATLAS FAX xrootd proxy will be upgraded to 3.3.3 during September.
NDGF			

Site	Status	Recent changes	Planned changes
	dCache 2.3 (Chimera) on core servers. Mix of 2.3 and 2.2 versions on pool nodes.		
NL-T1	dCache 2.2.7 (Chimera) (SURFsara), DPM 1.8.6 (NIKHEF)		
PIC	dCache head nodes (Chimera) and doors at 1.9.12-23 xrootd 3.3.1-1		Next upgrade to 2.2 in 16th September
RAL	CASTOR 2.1.12-10 2.1.13-9 (tape servers) SRM 2.11-1		
TRIUMF	dCache 2.2.13(chimera), pool/door 2.2.10		

FTS deployment

Site	Version	Recent changes	Planned changes
CERN	2.2.8 - transfer-fts-3.7.12-1		
ASGC	2.2.8 - transfer-fts-3.7.12-1		
BNL	2.2.8 - transfer-fts-3.7.10-1	None	None
CNAF	2.2.8 - transfer-fts-3.7.12-1		
FNAL	2.2.8 - transfer-fts-3.7.12-1		
IN2P3	2.2.8 - transfer-fts-3.7.12-1		
KIT	2.2.8 - transfer-fts-3.7.12-1		
NDGF	2.2.8 - transfer-fts-3.7.12-1		
NL-T1	2.2.8 - transfer-fts-3.7.12-1		
PIC	2.2.8 - transfer-fts-3.7.12-1	None	None
RAL	2.2.8 - transfer-fts-3.7.12-1		
TRIUMF	2.2.8 - transfer-fts-3.7.12-1		

LFC deployment

Site	Version	OS, distribution	Backend	WLCG VOs	Upgrade plans
BNL	1.8.3.1-1 for T1 and US T2s	SL6, gLite	ORACLE 11gR2	ATLAS	None
CERN	1.8.6-1	SLC6, EMI2	Oracle 11	ATLAS, LHCb, OPS, ATLAS Xroot federations	

Other site news

Data management provider news

Experiment operations review and plans

ALICE

- significant and efficient grid usage over the summer break
- CVMFS

- ◆ deployment campaign: 33 tickets closed, 19 open
- ◆ 1 small site (60 job slots) running it in production since 2 weeks, with good job success rates
- ◆ next: switch a big T2
- CERN
 - ◆ job submissions by ALICE and other VOs failing for many hours due to Argus incidents on July 21-22 (GGUS:95914 [↗](#)) and Aug 23 (INC:365019 [↗](#))
- KIT
 - ◆ 23 production files lost due to a damaged tape; as they already had 2 other replicas, only catalog entries needed to get adjusted
 - ◆ local SE access still unstable, but with reduced impact since the firewall upgrade on July 24-25; jobs cap at 6k since Aug 12
- RAL
 - ◆ on Fri Aug 23 many WN ran out of memory due to one ALICE user's broken jobs that allocated huge amounts of memory in a very short time
 - ◆ to stop the issue, ALICE jobs were banned for the weekend
 - ◆ the guilty tasks were removed and their owner has been admonished

It is fair to expect that most ALICE sites will migrate to CVMFS well before the official target date of April 2014.

ATLAS

- Distributed Computing Operations as usual in the past month. To be noted the ongoing efforts invested in the FTS3 commissioning and integration by ADC and DDM Ops, more details in the FTS3 task force report.
- as a reminder ATLAS reports the "daily" issues at the WLCG "daily", all reports available here: <https://twiki.cern.ch/twiki/bin/viewauth/Atlas/ADCOperationsDailyReports2013>

CMS

- Dirk Hufnagel joined the machine feature task force for CMS
- HammerCloud
 - ◆ HC jobs stopped to be run on August 21, when the certificate of the glideIn WMS frontend expired, fixed on August 28
- Argus calls timing out: INC:365019 [↗](#)
 - ◆ August 23: two CERN hosts we use for CMS analysis jobs submissions became unusable to users because gsissh connection was failing. We also had about 5K CMS pilots dying (condor Held) due to glexec failures. Problem apparently disappeared around noon, w/o any action taken
 - ◆ August 23: ALICE job submission was also badly affected from ~08:30 until ~14:00, causing CERN to get almost fully drained of ALICE jobs. What is more, also beforehand and afterwards there have been sporadic AuthZ errors, so it seems the Argus machinery was/is flaky
 - ◆ This issue seems to be to be created by high user activity, and the way CERN setup the HA Argus currently is not able to cope with this load. The system worked as designed in the sense that the overloaded servers were taken out of the alias and the load moved on to the next while the previous one was draining.
- T1 prioritization changes:
 - ◆ Following changes have been implemented recently:
 - ◊ t1production role is not used anymore after the GlideIn WMS frontends were merged, only the production role is used from now on

- prioritization is done on the frontend level => fractional shares within the CMS T1 allocations are not needed anymore, simple prioritization is sufficient
 - ◊ disk/tape separation allows to open the T1 CPU resources for analysis accessing samples on the disk endpoints
- ◆ Updated prioritization policy: simple prioritization without shares in CMS CPU allocation at the sites
 - ◊ available CPU resources should be assigned first to production role work, then to pilot role analysis work, then to any other role or no role analysis work within the CMS VO
 - ◊ Access for analysis to T1 CPU resources should prefer glxexec and therefore the pilot role
- ◆ All documented here
<https://twiki.cern.ch/twiki/bin/view/CMSPublic/CompOpsPoliciesVOMSRoles>
- CVMFS: working with the last remaining ~10 T1/T2 sites to switch to CVMFS
 - ◆ Seeing from time to time black holes due to CVMFS, instructed site admins to prepare debug tarballs and send them to CVMFS development team
- glxexec: after the summer break, we will restart working with the sites to enable glxexec
 - ◆ by late fall, we would like to make the CMS glxexec SAM tests critical and rely on glxexec in our global glideIn WMS pool
- Savannah->GGUS:
 - ◆ work continues to migrate all ticket activities to GGUS, working with the GGUS development team on:
 - ◊ new web input mask, test instance and implementation of additional CMS internal support units
 - ◆ list of open tickets: SAV:138640 [↗](#), SAV:134411 [↗](#), SAV:134413 [↗](#), SAV:134415 [↗](#), SAV:134416 [↗](#)

Maria D. adds that discussions on the migration to "pure" GGUS for CMS tickets will resume next week with a meeting with the GGUS developers. Another meeting has been scheduled in the second week of October.

Helge adds that the Argus problem was basically due to an overload and additional servers will be deployed next week.

LHCb

- FTS3 test OK
- python bindings for grid middleware nearly ready for python 2.7
- incremental stripping campaign will be starting mid September for 2 months
- GridKA : tape family set for the next round of data taking in order to improve the tape access
 - ◆ Info for T1 sites about tape families: <https://twiki.cern.ch/twiki/bin/view/LHCb/NameSpace>

FTS 3 was used to run transfers to/from CERN. The only problem was with the accounting, due to a wrong timestamp published for the transfer start times. This only happens when using the old gLite client. Next week a new version of the server will fix the issue.

Task Force reports

CVMFS

- Very good progress for ALICE site migration to CVMFS (campaign started in July)
 - ◆ 31 of 51 sites have deployed CVMFS, most of the remaining sites promise to deploy this fall

- ◆ 2 sites have not acknowledged the ticket yet (ITEP, IN-DAE-VECC-02)
- New GGUS support unit "CVMFS" added under new category "File System"

SHA-2 migration

- EGI Operations Management Board Aug 27 meeting [☞](#)
 - ◆ good progress for CREAM, VOMS, WMS
 - ◆ little progress for StoRM, dCache
 - ◇ no SHA-2 or JGlobus-2 mentioned yet in the dCache 2.2 series release notes [☞](#) (the 2.6 series is OK)
 - ◆ discussion outcomes:
 - ◇ aim for **Dec 1** deadline for compliance of services
 - ◇ ask IGTF/EuGridPMA to delay SHA-2 timeline by 2 months
- CERN's VOMS servers **can** generate SHA-2 proxies!
 - ◆ but SHA-2 certs need to be registered as **secondary** certs for now, as described here
 - ◆ migration from VOMRS to the SHA-2 compliant VOMS-Admin is foreseen to happen in the autumn
- ATLAS and LHCb SW looks ready for SHA-2
- CMS
 - ◆ DBS2 will **not** be able to deal with SHA-2
 - ◆ it is scheduled to be retired in Nov
- ALICE: various SW still to be checked

Burt reports that FNAL is very busy implementing the disk/tape separation of its storage and he thinks that they should be able to upgrade dCache to the new SHA-2 compliant golden release by the new extended deadline - but it will not be easy.

Michel proposes to closely keep track of progress and issues and Burt agrees to report for the next meeting.

Massimo asks if Grid clients are concerned in the SHA-2 compliance. Maarten says that very old clients might have some issues in principle, but it looks quite unlikely that we would run into trouble there; no issues with clients (oldish and new) were observed in relation with SHA-2.

For the next meeting, Maarten will give more information about the SHA-2 testing in ALICE.

Nicolò asks Peter to communicate the exact software version numbers used in EGI to assess SHA-2 compliance; they should be compared to the WLCG baseline versions.

Stefan and Simone clarify that LHCb and ATLAS tested their central services and frontends for SHA-2, but did not attempt to run jobs and would like to know if it is useful. Andrea S. comments that CMS tested also the full job submission chain. Maria G. recommends all experiments do the same. Maarten says that a small set of such tests should be enough and that the matter can be discussed further offline.

It is clarified that the migration of user records from VOMRS to VOMS-Admin will be done automatically.

SL6

- **Tier1**
 - ◆ Done: 8/16 (Alice 4/10, Atlas 6/13, CMS 3/9, LHCb 4/9)
 - ◆ In progress: CCIN2P3 and NDGF (both as planned)
 - ◆ TW is upgrading and already in production in Atlas with 50% of the resources. The other half will be done by the end of the week.
 - ◆ Others: there are 5 Tier1s which put as a deadline 2013-10-31. They might be able to do it sooner (for example Kit and CNAF commented on working towards September). RAL is

deciding if to go down the ARC-CE/condor route or if to keep Torque/maui. In the first case it will be just matter of rolling out WNs in the second case we need to organise a big bang upgrade. RRC-KI-T1 and FNAL haven't finalised their plans.

- ◆ Tier1s status: <https://twiki.cern.ch/twiki/bin/view/LCG/SL6DeploymentSites#Tier1s>

- **Tier2**

- ◆ Done: 48/129 (Alice 11/39, Atlas 28/89, CMS 22/65, LHCb 13/45)
- ◆ A number of sites that put as deadlines 2013-07-31 and 2013-08-31 or marked themselves in progress are actually postponing.
- ◆ 7 sites are still marked with "no plans and no testing" I will start to open tickets for these sites starting from 1st of September.
- ◆ Few sites are waiting for Quattor templates. Investigating when these templates will be available.
- ◆ Tier2s status: <https://twiki.cern.ch/twiki/bin/view/LCG/SL6DeploymentSites#Tier2s>

- **EMI-3**

- ◆ voms-clients have been fixed and the latest version is in the PT repository but not in EMI-3 yet. <http://italiangrid.github.io/voms/2013/08/21/voms-clients-3-0-3-released.html>
- ◆ With this latest patch both Atlas and CMS run fine. CMS commented that they can run on DPM because they use xrootd but since this is a requirement for CMS there shouldn't be any problem.
- ◆ 5 sites, supporting LHCb, Atlas and CMS, are currently in production with SL6/EMI-3 combination: DESY-HH, UKI-LT2-Brunel, UKI-NORTHGRID-LIV-HEP, ASGC, wuppertal.
- ◆ EMI-3 status:
https://twiki.cern.ch/twiki/bin/view/LCG/SL6DeploymentSites#EMI_3_Experiments_Storage_matrix

- **HS06**

- ◆ Reminder that sites are requested to run HS06 benchmark and update the value in the BDII. Sites which have already done it see increase up to 25% depending on the hardware. It would help the community if sites added their results to the hepix site http://w3.hepik.org/benchmarks/doku.php?id=bench:results_sl6_x86_64_gcc_445 however it is not clear how sites can have access to it. The normal hepik login doesn't work. Investigating with Hepix people.

It is not fully understood why re-running the HS06 benchmark on SL6 gives higher results compared to SL5. It might be due to the new compiler or the new kernel. Considering that most applications (e.g. for CMS) are still compiled on SL5, Ian thinks that the binaries compiled with the SL5 system compiler should be used as reference. The effect of the compiler should be properly measured.

Helge suggests to discuss the increase in the benchmark values also in the WLCG MB.

Tracking Tools

Brainstorming is starting now about how to implement the possibility to notify multiple sites via a single GGUS ticket. You are invited to contribute comments in Savannah:138299.

This is a complex development so we have to be clear about the requirements. We also need to understand how often this bulk submission is likely to be needed. Summarising what we have so far:

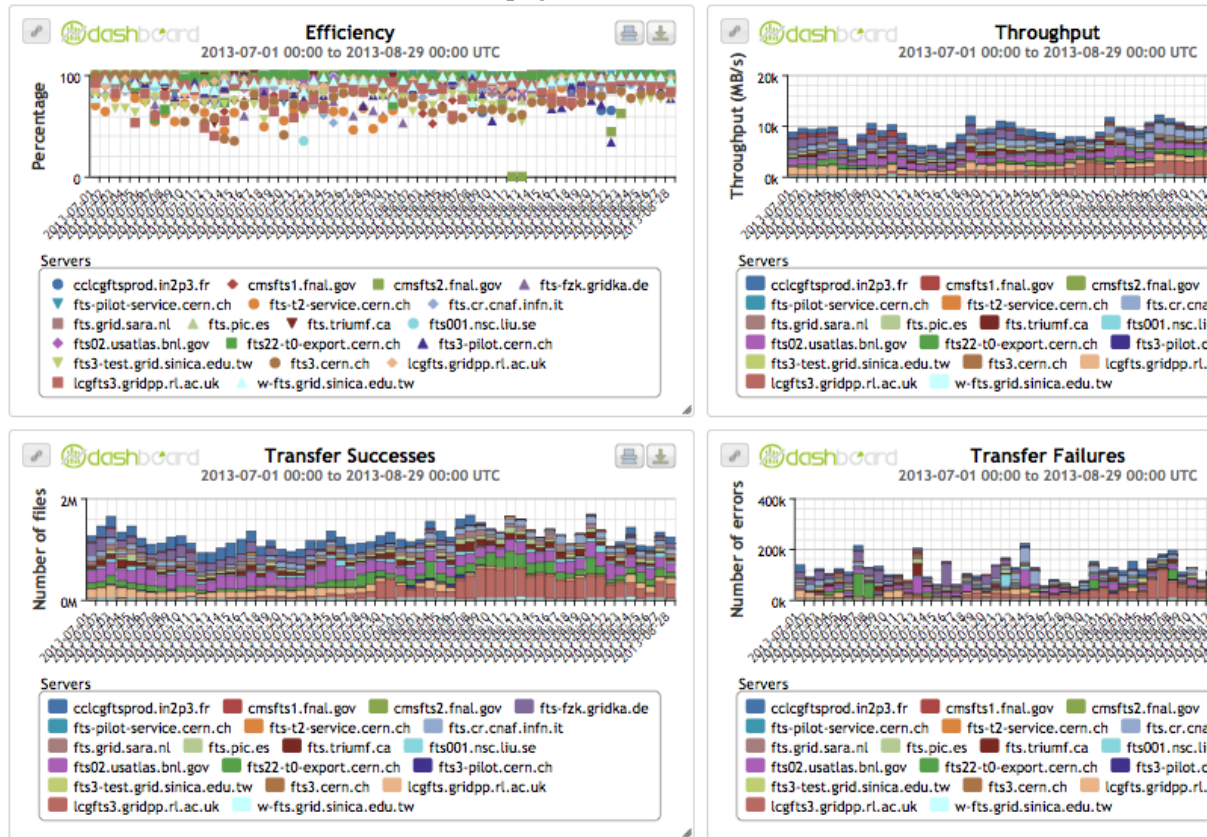
1. To avoid mistakes and abuse, the privilege should be offered to expert supporters/submitters.
2. The submitter name should remain their own (no TPM involvement) so they control the future exchanges.
3. Subject and Description field should be typed once and cloned to all individual ticket instances.
4. The official site name, as in GOCDB or OIM, should still be offered for selection to the expert submitter, so they don't risk to mis-type complex names, e.g. AEGIS03-ELEF-LEDA (maybe not a WLCG name but the prettiest I found).

5. The possibility to select >1 ROC/NGI (not only "Notify Site") should also be foreseen.

FTS-3

- ATLAS, CMS, LHCb are now all using FTS3 for functional tests or real production.
- FTS3 is used now in ATLAS for all production activity for approximately 30% of the overall transfers.

◆ Screen_Shot_2013-08-29_at_4.22.57_PM.png:



, WLCG dashboard 1st July - today

- The integration with the experiment is an iterative process (O(10) patches in the last 2 months)
- we had experienced problems (starting from around 15th of August). The first problem we observed was the one of not having enough ACTIVE transfers in some links (link plenty of SUBMITTED).
 - ◆ this has been understood and solved, FTS3 DB was not optimized properly.
 - ◆ a patch was applied the 22nd of August, but this update had side effects which lead into having links completely stuck, some of the jobs were not reporting back their statuses.
 - ◆ we have understood the problems in the past days, and just yesterday (28th) a new FTS3 update has been releases and deployed in RAL.
 - ◇ we are evaluating the performances right now, it seems that the amount of errors is much higher than before. it seems that there are links with very low throughput which have an impressively high number of ACTIVE O(100)
 - ◇ since this morning we agreed with FTS3 devs to start putting more load with ATLAS functional tests on the fts3.cern.ch production CERN endpoint which is running with MySQL backend (i.e. similar to RAL), to try to reproduce - and then solve - the issues.
- a status report provided the 26th of August in an email sent to the fts3-steering mailing list
<https://groups.cern.ch/group/fts3-steering/Lists/Archive/Flat.aspx?RootFolder=%2Fgroup%2Ffts3-steering%2F>

Alessandro will put the status report also in the task force twiki.

XrootD Deployment

- New version of GLED Collector for XRootD detailed monitor (M. Tadel)
 - ◆ Change log relative to 1.4.0:
 - ◇ external updates:
 - ◇ root 5.34.07 -> 5.34.09
 - ◇ apr 1.4.6 -> 1.4.8
 - ◇ activemq-cpp 3.7.0 -> 3.7.1
 - ◇ resolve clients with numeric IPs (needed as more and more servers are getting configured to not do the reverse lookup)
 - ◆ RPMs available in the WLCG repo.
- dCache xrootd door monitoring plugin
 - ◆ new version with improved functionality will be made available in the WLCG repo by next week (announcement will be circulated)
 - ◆ In the meantime, few ATLAS sites (MWT2 and AGLT2) have already deployed the new version.

gLExec

- 39 tickets closed and verified, 55 still open, many on hold until mid/late autumn (in line with SL6 migration)
- Deployment tracking page
 - ◆ various countries all done, in particular France!

AOB

Action list

1. Tracking tools TF members who own savannah projects to list them and submit them to the savannah and jira developers if they wish to migrate them to jira. AndreaV and MariaD to report on their experience from the migration of their own savannah trackers.
2. Investigate how to separate Disk and Tape services in GOCDB
 - ◆ proposal submitted via GGUS:93966 [🔗](#)
3. Agree with IT-PES on a clear timeline to migrate OPS and the LHC VOs from VOMRS to VOMS-Admin
 - ◆ **in progress**
4. Experiments interested in using WebDAV should contact ATLAS to organise a common discussion
 - ◆ **input welcome by the next GDB**
5. Contact the storage system developers to find out which are the default/recommended ports for WebDAV
 - ◆ **in progress** Recommendations for WebDAV deployment

-- AndreaSciaba - 28-Aug-2013

This topic: LCG > WLCGOpsMinutes130829

Topic revision: r40 - 2013-08-30 - MaartenLitmaath



Copyright &© 2008-2024 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback