

CernVM, Systems & Services Plan Of Work for 2020

G Ganis

On behalf of the CSS team

10 February 2020

Outline

- Quick reminder about the CSS team
 - Team members
- System & Services
 - Last year activities
 - A plan of work for 2020
- CernVM plan of work
 - Presented on 20 January 2020 by [J Blomer](#)

S&S PoW 2019 presentation

- G Ganis, [System & Services Plans for 2019](#), 21 Jan 2019

CSS Team mandate

Create the conditions and synergies to:

- Develop a *Software Provisioning Service*
 - From the tool for efficient software distribution to the building and validation of the complete software stacks
- Develop services and/or products to support the software development activity

CSS Manpower evolution

		2019	2020	2021	
G Ganis	STAF	60%	40%	40%	S
J Blomer	STAF	50%	50%	50%	C
I Goulas	STAF	100%	100%	100%	S
S M Muzaffar	STAF	10%	10%	10%	S
S Mosciatti	TECH,FELL	60%	100%	66%	C
R Popescu	FELL	66%			C
A Stano	FELL	80%	100%	20%	W
J Heinz	TECH	80%			S
R Bachmann (+TBS)	TECH	33%	80%+	80%+	S
J Priessnitz (+TBS)	TECH	60%	40%+	50%+	C
I Razumov	PJAS	100%	100%	40%	G, S
Russian WLCG	USER	100%	100%	100%	G, S
Total (FTE)		8.1	7.2+	5.6+	

C: CernVM
S: Systems & Services
G: Genser/GeantVal
W: Web



CMS librarian
hosted by SFT

Russian WLCG

- D Konstantinov
- G Latyshev
- A Schukin

Summer Stud. 2019

- M Marik (S)
- S Isgandarli (C)
- M Mokrzan (C)

Valuable contribution
from P Mato

Team members activity reports in 2019

CernVM

- S Mosciatti, [A report on unpacked, DUCC and ingest](#), 1 Apr
- J Blomer, [The new CernVM-FS release](#), 15 Apr
- J Blomer, [CernVM Users Workshop Report](#), 24 Jun
- J Priessnitz, [Performance engineering of S3 storage backend](#), 9 Dec

SPI

- S Muzaffar, [Jenkins service status and evolution](#), 30 Sep
- I Goulas, [The JIRA SFT instance](#), 30 Sep
- J Heinz, [Technical Student Report](#), 28 Oct
- I Razumov, [Spack Plans in SPI](#), HSF Packaging Group Meeting, 11 Dec

Summer Student reports

- S Isgandarli, [Deep I/O performance analysis using modern Linux tools](#), 12 Aug
- M Marik, [Porting the LCG software stack to the ARM architecture](#), 9 Sep
- M Mokrzan, [A JavaScript based CernVM-FS monitor](#), 16 Sep

CHEP 2019 contributions

- SPI: [Optimizing Provisioning of LCG Software Stacks with Kubernetes](#)
- CernVM-FS: [A fully unprivileged CernVM-FS](#)
- R&D: [Towards a Turnkey Software Stack for HEP experiments](#)

System & Services (aka SPI)

S&S deliverables

- Software stack of **about 450 packages**
- **Large number of combination**
 - {OS1, OS2, ...} x {Comp1, Comp2, ...} x {opt, dbg} x Python{2,3}
- **Nightlies**
 - C: /cvmfs/sft-nightlies.cern.ch
- **Releases**
 - C: /cvmfs/sft.cern.ch
 - R: RPMs
 - T: tarballs
 - D: Docker images (experimental)

Snapshot of the current situation

	SLC6	CentOS7	Ub18 LTS	MacOsX	
Native			T, C (D)	(T) (C)	
gcc 6.2, 7	T, C, R (D)	T, C, R (D)			legacy
gcc 8	T, C, R (D)	T, C, R (D)	T, C (D)		pro
gcc 9		T, C, R (D)			dev
clang 8		T, C, R (D)			pro
clang 9		T, C, R (D)			dev

() = could be provided

S&S infrastructure and Web Server

Build infrastructure:

- Jenkins-based orchestration
- Openstack-provided VMs
 - SLC6, CentOS7, Ubuntu LTS, Ubuntu latest, Supported Fedora
 - CentOS7 docker-host: slc6, centos7, ubuntu+, fedora containers
 - Physical machines for MacOSX, ARM, GPU-enabled, ...
- Staging/shared area on EOS

Web Server for RPMs and tarfiles: EOS

- <https://lcgpackages.web.cern.ch/lcgpackages/>
 - Binary TAR balls:
 - tarFiles, tarFiles/sources
 - YUM repo:
 - rpms_contrib, rpms, rpms_updates
 - YUM repo configuration files at [rpms/etc/yum.repos.d/lcg.repo](https://rpms.web.cern.ch/rpms/etc/yum.repos.d/lcg.repo)

S&S stakeholders

	ATLAS	LHCb	FCC, SWAN	BE	CLIC, NA61, NA62, other SME, ...
Nightlies (C)	✓	✓	✓	✓ (dedicated views)	✓ (selections)
Releases (C)	✓		✓		✓ (selections)
Releases (R)	✓	✓			

- Many users on lxplus and elsewhere
 - Increasing over time due to the phase out of /afs/cern.ch/sw

Activities in 2019

Plan Of Work for 2019

- Consolidation
 - Infrastructure (Jenkins Server, Nodes), Containers
- Improvements
 - CernVM-FS publication, RPM repository, Testing
- Review of package content
 - Review of needs for debug symbols
- Follow requests for new architecture / platforms
 - ARM, Power, ...
- Documentation
- Drupal 8 move
- AoB

Last year activities: highlights (1)

- Infrastructure

- New, puppetized, Jenkins server
- Consolidated container infrastructure
 - Definition and build of images, ...
 - New Ubuntu-based docker host for latest Linux distributions
- Investigations of Kubernetes
- Summary of the builds: [dashboard](#)

- LCGCMake

- Consolidation of bin/lcgcmake
- *Layered Stack* concept: design and implementation

- Release procedure, build and deployment

- Consolidation of the Release Candidate concept
- Optimization of CernVM-FS catalogue branching
- Better partitioning of RPM repository

Nightly build report for 2020-02-08 (generated at 10:25)

Clean

dev3

Platform	Started	Build			Copy		Post-build		Install		Test			Completed	isDone file timestamp
		Node	Duration	Result+Log	Duration	Status+Log	Status	Duration	Status+Log	logtest	import	shell			
x86_64-centos7-ci4ngli-opt	02:06:02	lgapp-centos7-x86-64-33	1:35:13	✓	0:00:13	✓	✓	0:13:02	✓	✓	✓	✗	✗	04:43:45	✓
x86_64-centos7-gcc62-opt	02:06:02	lgapp-centos7-x86-64-42	1:50:08	✓	0:00:11	✓	✓	0:13:00	✓	✓	✓	✗	✗	05:09:57	✓
x86_64-centos7-gcc8-dbg	02:06:02	lgapp-centos7-x86-64-40	1:27:22	✓	0:00:16	✓	✓	0:15:21	✓	✓	✓	✗	✗	04:30:38	✓
x86_64-centos7-gcc8-opt	02:06:02	lgapp-centos7-x86-64-39	1:53:28	✓	0:00:40	✓	✓	0:12:55	✓	✓	✓	✗	✗	04:56:53	✓
x86_64-centos7-gcc8-dbg	02:06:02	lgapp-centos7-x86-64-38	0:00:15	✗		✗	✗		✗	✗	✗	✗	✗		✗
x86_64-centos7-gcc8-opt	02:06:02	lgapp-centos7-x86-64-74	2:06:30	✓	0:00:11	✓	✓	0:13:16	✓	✓	✓	✗	✗	05:23:17	✓
x86_64-ic8-gcc8-dbg	02:06:02	lgapp-centos7-x86-64-40	0:00:15	✗		✗	✗		✗	✗	✗	✗	✗		✗
x86_64-ic8-gcc8-opt	02:06:02	lgapp-centos7-x86-64-43	0:00:15	✗		✗	✗		✗	✗	✗	✗	✗		✗

dev3python3

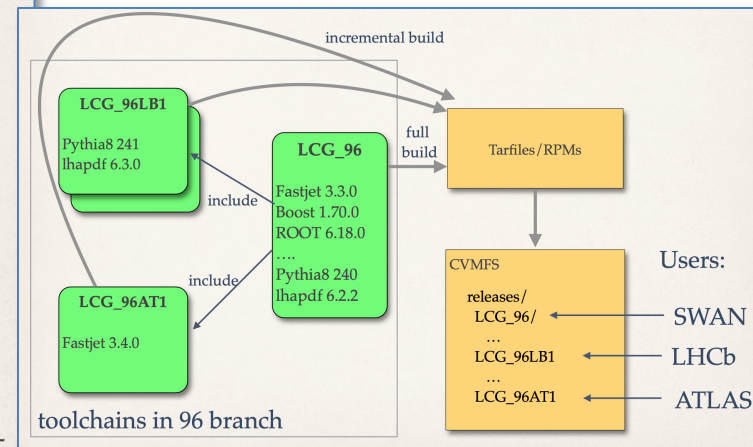
Layered Stack concept

Problem statement

- * LCG stacks with **multiple versions** of some packages can be problematic
 - * Difficult to configure the exact versions of dependencies to use
 - * Cannot easily have two configurations of a given package with different set of dependencies
- * Adding **new versions** of packages (mainly in MC generators) to an existing LCG release is also problematic
 - * E.g. adding a new version of LHAPDF can break existing releases

Proposed solution

- * Ensure that ALL provided LCG stacks are always consisting of a **single version of each package**
 - * Avoid the the current practice that we have several versions
- * Be able to **build stacks on top of existing stacks** in an incremental manner
 - * Package versions can be over-written
 - * In same cases packages could be removed from the stack
 - * The version change of a package may trigger the build of all dependent packages



(from [P Mato @ LIM-23/5/2019](#))

Last year activities: highlights (2)

- **Testing**
 - Introduced *single Python package* and *environment setting* testing
 - Prototyped running of roottest
- **Package content**
 - Removed GRID packages
 - Which were an outdated repackaging of EPEL RPMs
 - Increased usage of system packages
- **New version of [lcginfo](#)** , including exact package content per platform
 - Revised version of lcgdocs
- **Support**
 - Consolidation of [JIRA](#) as **main support channel and planning tool**
 - Massive reduction of open JIRA tickets
 - Service Now Functional Element: Software Development For Experiments
 - Low activity (1-2 /month) mostly connected to AFS phase out
- **New packages of general interest**
 - Gaudi, PODIO, ACTS-core, ...
- **Many new generators / versions included**

LCG_94a → LCG_96b
+ 84 new packages
- 22 removed

New lcginfo: lcginfo.cern.ch

LCG Info: Release 96b

[Home](#) / [Releases](#) / 96b

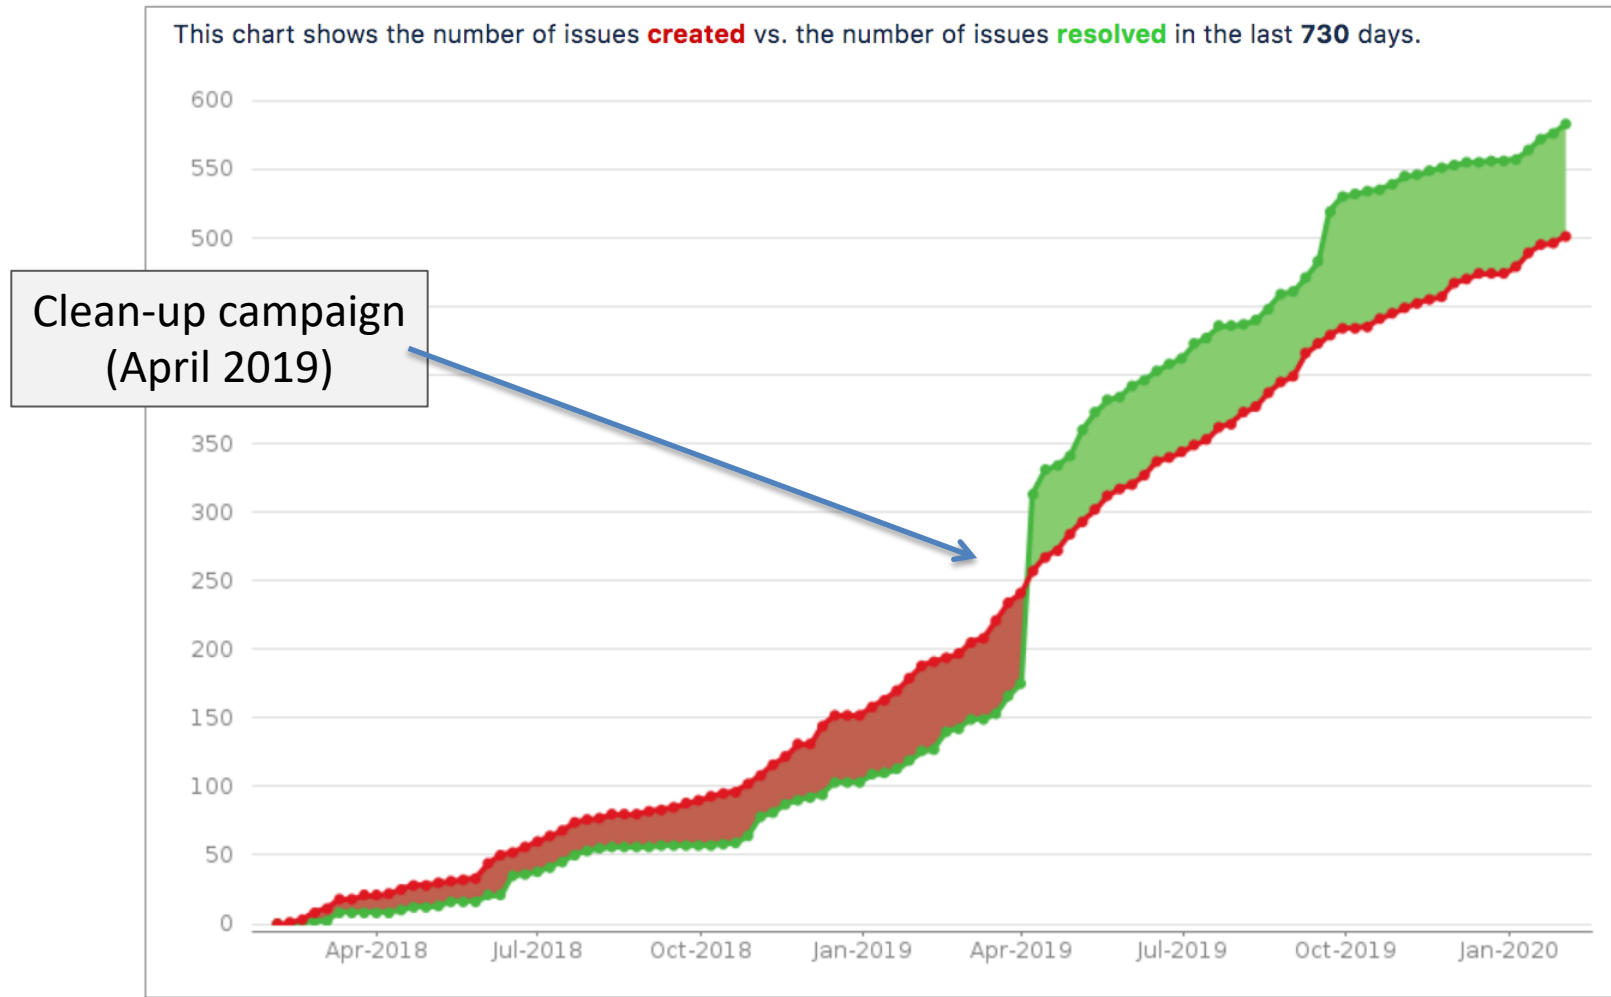
Compare to another release:

Description: Publication of patch release LCG_96b

Release date: Sep 25, 2019

Platforms:		Packages	Release Notes
	x86_64-centos7-clang8-dbg	Packages	Release Notes
	x86_64-centos7-clang8-opt	Packages	Release Notes
	x86_64-centos7-gcc62-opt	Packages	Release Notes
	x86_64-centos7-gcc8-dbg	Packages	Release Notes
	x86_64-centos7-gcc8-opt	Packages	Release Notes
	x86_64-centos7-gcc9-dbg	Packages	Release Notes
	x86_64-centos7-gcc9-opt	Packages	Release Notes
	x86_64-slc6-gcc8-opt	Packages	Release Notes
	x86_64-ubuntu1804-gcc8-dbg	Packages	Release Notes
	x86_64-ubuntu1804-gcc8-opt	Packages	Release Notes

JIRA created/resolved tickets



Last year activities: highlights (3)

- On-demand processing of Merge Requests in GitLab
 - Can test MR for dev4, dev3, dev3python3 on the relevant platforms
- First look at Spack for the LGC stack
 - Presented at LIM and [HSF](#)
 - LCG_96 equivalent binaries and recipes available, and [documented](#)
- New Web Site (next slide)
 - Single entry for all relevant information
 - LIM announcements and minutes, JIRA, Documentation, ...
- Drupal 8 (or other) migration
 - Completed or in production
 - EP Department, EP R&D, Neutrino-Secretariat, User-Office
 - Work Started
 - EP-SFT, EP-ESE, CernVM (Jekyll)
- AFS area [/afs/cern.ch/sw/lcg](#) made inaccessible
 - Increased requests / users of [/cvmfs/sft.cern.ch/lcg](#)

SPI web site: spi.web.cern.ch

SOFTWARE PROCESS INTEGRATION



Software Stacks For CERN Experiments

In the context of the EP-SFT activities, the SPI team provides software stacks for users in general. These software stacks consist of several layers, including Python modules - all available for a large number of platforms.

SPI SERVICES



DOCUMENTATION

USEFUL LINKS

PUBLICATIONS

LCG Releases

Software packages like mathematical libraries, DBs, Compilers etc which used in specific software at CERN are recompiled and made available for specific platforms.

LIM

The Librarians and Integrators Meeting (LIM) is a forum bringing together librarians of all the LHC and other experiments to share experiences about software stack provisioning. The forum meets regularly - usually every two weeks - to discuss issues about the builds, new compilers or package versions, new operating system versions, et etc.

[LIM Meetings minutes](#)

Drupal 8 migration revisited

- CERN strategy for Drupal 7->8 migration was modified to cope with complexity of the move
- **Drupal 8 remains CERN's recommended CMS**
 - Migration support either from IT or external company(ies)
 - Review process to identify sites staying-on/leaving Drupal 8
- **Opening to alternative solutions** for sites not strictly requiring Drupal functionality

OpenShift	Container based hosting solution		Production	
MicroServices	TBD, use-case based (e.g. <u>WordPress, Jekyll, Hugo, etc.</u>)			Web Frameworks
Drupal	For CMS features		Production	
Elasticsearch	Tailored search service / Elasticsearch / Invenio		Production	

[From MALT@CERN](#)

Web sites under SFT responsibility

- Review process suggested to go for an external company for {EP, EP News, CernVM} web sites
- After internal discussion **decided to opt out from price inquiry** and continue the **migration in-house as planned**
- ROOT and CernVM have opted for an alternative technology
 - Jekyll

Last year activities: releases

Releases

- LCG_95{a}, python3: ROOT v6.16/xx
- LCG_96{b}, python3: ROOT v6.18/xx
 - LCG_96c_LS: test of Layered Stack concept (see later)

New views

- For BE
 - devBE: mostly python packages
 - devNXCALS: dedicated version of dev3 for usage in SWAN
- dev{3,4}cuda
 - Build with CUDA support for TensorFlow, ...
- devARM
 - Port on ARM (tested by LHCb and ATLAS)

Plan Of Work for 2020

- Infrastructure consolidation
 - Jenkins Server Migration, Nodes setup, Containers
- Improvements / Developments
 - LCGCMake, CernVM-FS publication, Testing
 - Spack
 - Kubernetes
- Support
 - Users, CentOS8 port, new tools / platforms, ...
- Web sites
 - Drupal 8 or other
- Documentation

Infrastructure

- Jenkins

- Finalize migration to new Server instances

[Q1]

- Complete ROOT, CernVM
 - Move {SPI, G4, FCC} jobs to new instance

- Nodes

- Move away HOME from AFS
 - Rationalization of Puppet branches
 - Ensure that *production* contains the latest updates

[Q1]

- Consolidate use of containers

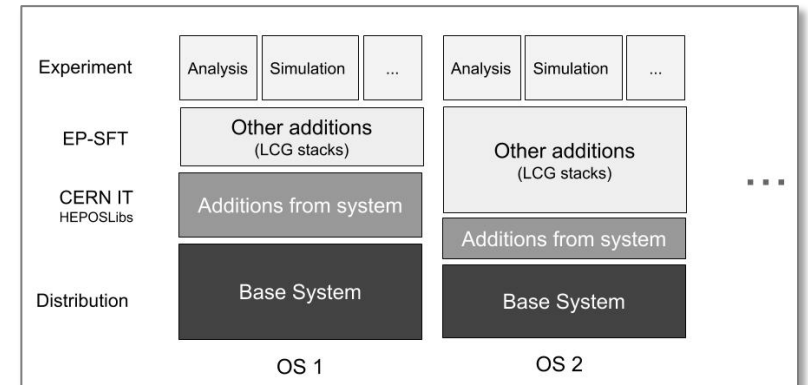
- Fully container-based system

- Releases still produced on VMs

[Q1/Q2]

Baseline System Definition

- **LCG stacks** stand between experiment needs and system provided components



- Target is to **critically review the package content** to **increase** what is taken from the system (HEP-OSLibs)
 - Look also at EPEL ?
- **In collaboration with IT and customers**
 - Dedicated LIM(s) to make concrete progress

Improvements: LCGCMake

- Single source for binaries
 - Use tarballs for CernVM-FS and RPMs
- Consolidate Layered Stack concept
- RPM technology
 - Single MetaRPM
 - Include RPM creation in LCGCMake
 - Now separate script
 - Improve RPM structure
 - Standard use OS, ARCH, versioning
 - Currently:
`ROOT-a1638_6.14.08_x86_64_centos7_gcc8_dbg-1.0.0-94.noarch.rpm`
 - Investigate *standardization* of RPM repository structure
e.g.
 - `lcgrepo/7/96/x86_64/Packages`
 - `lcgrepo/7/96/x86_64/debug/Packages`

[Q1/Q2]

Improvements: CernVM-FS publication

- Use of CernVM-FS Conveyor high-level interface
 - See [CernVM-FS high-level publication interface](#)
 - Repository Gateway, RabbitMQ, Job DB
 - Allows to optimize operations, handling dependencies and maximizing concurrency
- Goal is to **liberate CI pipeline from handling cvmfs_server intrinsics**
- **Prototype available**, estimate 1-2 months to develop a first usable version
 - In collaboration with CernVM core team
- Plan is to use **new, HSF oriented, repositories**
 - sw.hsf.org, sw-nightlies.hsf.org
 - CernVM-FS Gateway-enabled

[Q1/Q2]

Improvements: testing

- Increase coverage for releases or nightlies
- Finalize inclusion of *roottest* in regular tests
 - Run other similar test suites?
- Investigate the possibility to run integration tests provided externally
 - E.g. from experiments

Developments

- **Spack**
 - Continue investigations / testing
 - New version expected by Q2
 - Summer student project
 - Coordinate with HSF Software Tools and Packaging WG
- **Containers**
 - Alpha version of a fully Kubernetes-based build service
 - RB technical student project

Support

- LCG stack customers
 - And users of /cvmfs/sft.cern.ch (as AFS replacement)
- CentOS8 port
- Follow requests for new architecture, platforms, compilers, tools
 - PowerPC, TAU, OneAPI, ...
 - In coordination with OpenLab
- Follow requests for new packages and versions

AoB


- Infrastructure managing tools
 - Bring latest improvements in production
 - GitLab sft/lcgjenkins repository
- Streamline production of release containers
 - Investigate use of CernVM-FS shrinkwrap to optimise the images

Web sites

- Migration plan to Drupal 8 or other
 - Department
 - To be completed
 - EP-SFT (w/ some revision), EP-ESE
 - Next
 - EP-AGS, EP-News, EP-DT (tbc)
 - Help migration of group specific sites
 - CernVM (Jekyll), Geant4 (D8), possibly ROOT (Jekyll)
- Support for existing D7 sites until migration done
 - EP-SFT, ...

Status and Plans presentation at EP-SFT group meeting on Monday 24 Feb 2020

Service Tasks

Service/Task	Main Responsible	Alternate	Documentation
Jenkins service	Shahzad Muzaffar	Gunter Folger	HowTo
Coverity service	Gabriele Cosmo	Axel Naumann	
CDash service	Benedikt Hegner	Guilherme Amadio	HowTo
Puppetized nodes	Shahzad Muzaffar	Gunter Folger	
Windows nodes	Bertrand Bellenot	Gunter Folger	
Mac nodes	Axel Naumann		
Other OS nodes	Gerardo Ganis	Axel Naumann	HowTo
Drupal Manager	Andrea Stano		
ITUM contact	Jakob Blomer	Graeme Stewart	
C5 contact	Jakob Blomer	Enric Tejedor	HowTo
Jira Service	Ilias Goulas		
Grafana	Ilias Goulas		HowTo 
Training	Enric Tejedor	Danilo Piparo	

Still few cases available: please volunteer!
Need an effort to provide documentation!

Outlook for 2020

- Consolidation of existing infrastructure
- Continue evaluation/investigation of new products
 - Spack, as package manager
 - Use of Kubernetes
- Towards a sustainable, for the group and the community, Software Provisioning Service