

CernVM, Systems & Services Plan Of Work for 2021

G Ganis

On behalf of the CSS team

25 January 2021

Outline

- Quick reminder about the CSS team
 - Team members
- System & Services
 - Last year activities
 - A plan of work for 2021

The [key4hep plan of work](#) will be presented in a separated talk today

The [CernVM plan of work](#) will be presented on 15 February 2021

For reference: [System & Services Plans for 2020](#)

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

		2020	2021	2022	
G Ganis	STAF	40%	20%	20%	S
J Blomer	STAF	50%	50%	50%	C
A Sailer (1)	STAF	34%	66%	66%	S, K
I Goulas	STAF	100%	100%	100%	S
S M Muzaffar	STAF	10%	10%	10%	S
S Mosciatti	FELL	100%	66%		C
A Stano	FELL	100%	10%		W
P Fernandez Declara (2)	FELL	80%	100%	20%	K
V Volkl	FELL	100%	100%		K
R Bachmann + A Kraut (+TBS)	TECH	100%	100%	100%	S
J Priessnitz + A Valenzuela (+TBS)	TECH	90%	100%	100%	C
I Razumov (+TBS)	PJAS	100%	100%	100%	G, S
Russian WLCG	USER	100%	100%	100%	G, S
Total (FTE)		10.0	9.2	6.8+	

C: CernVM
 S: Systems & Services
 G: Genser/GeantVal
 W: Web, K: Key4hep

 CMS librarian
 EP-R&D

Russian WLCG
 - D Konstantinov (1)
 - G Latyshev

TECHs who left in 2020
 - R Bachmann (S)
 - J Priessnitz (C)

Summer Students
 - M Tyagi (C), GSoC

Valuable contribution
 from P Mato

(1) LIM chairs
 (2) CREMLIN+ till 2/21

Team members activity reports in 2020

CernVM

- J Blomer, [Software Deployment pre-GDB summary](#), 11 May
- J Priessnitz, [Final studentship report](#), 22 Jun
- S Mosciatti, [Update on the integration between CVMFS and container technologies](#), 14 Sep

SPI

- A Stano, [Status of Drupal 8 migration](#), 24 Feb
- D Konstantinov, [Optimization of Pythia 8](#), 2 Mar
- R Bachmann, [Kubernetes for SPI nightly builds](#), 18 May
- R Bachmann, [Technical Studentship report](#), 3 Aug

AoB

- G Ganis, J Blomer, [IT Matters](#), 30 Mar
- J Blomer, [Summary of the ITUM meeting](#), 13 Jul
- J Blomer, [Summary of the ITUM meeting](#), 26 Oct

System & Services (aka SPI)

S&S deliverables

- Software stack of **about 500 packages**
- **Large number of combinations**
 - {OS1, OS2, ...} x {Comp1, Comp2, ...} x {opt, dbg} x Python{3,2}
- **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 (last release LCG_99)

	CentOS7	Ub20 LTS	MacOsX	CentOS8
gcc 8	T, C, R, (D)			
gcc 9	T, C, R, (D)	T, C, (D)		
gcc 10				T, C, R, (D)
clang 8	T, C, R, (D)			
clang 10	T, C, R, (D)			
clang 12			T, C	

 Native

() = could be provided

S&S infrastructure and Web Server

Build infrastructure:

- Jenkins-based orchestration
- Openstack-provided VMs
 - IT supported flavours: CentOS7, CentOS8
 - CentOS7 and CentOS8 docker-host:
 - Ready to use docker images: slc6, centos7, ubuntu{16,18,20}, fedora, ...
 - Physical machines for MacOSX, ARM, GPU-enabled, ...

EOS Web Server for RPMs and tarfiles: <https://lcgpackages.web.cern.ch/>

- Staging/shared area
- Binary TAR balls: tarFiles, tarFiles/sources
- RPM repo: [lcg/repo/{6,7,8}/LCG_NN](#), [lcg/contrib/{6,7,8}](#)
 - Legacy repositories: rpms_contrib, rpms, rpms_updates
 - YUM repo configuration files at [lcg/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)	✓	✓			

- Partial stacks provided for development builds to specific projects
 - E.g. Geant4, VecGeom, DD4hep, AdePT, ...
- Many users on lxplus and elsewhere
 - Increasing over time due to the phase out of /afs/cern.ch/sw

Activities in 2020

Impact of sanitary crisis

- Productivity remained high
 - People managed to work rather well, also remotely
- Main effect was to cancel the summer student project
 - Ready to do remote supervision, but selected candidate found alternative occupation
 - The confirmation of the program came too late

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**

Last year activities: highlights (1)

- **Infrastructure**

- Jenkins: single monolithic server replaced by new puppetized servers
 - 3 instances deployed on large VMs: ROOT, CermVM, SPI/G4
 - Features: quick and regular updates, single sign-on
- Puppetized AFS-free CentOS{7,8} docker host build nodes commissioned
 - Ubuntu-based docker hosts replaced by CentOS8 ones
 - Newer kernel able to serve well all Linux flavours
- Consolidated, Jenkins-driven, infrastructure for building container images
- Puppet: rationalization of manifests
 - [qa](#) sync'ed into [production](#): back to normal workflow

- **Nightlies/Release procedure, build and deployment**

- Fully based on containers (no direct use of VMs any more)
- Moved to use Jenkins pipelines

Jenkins pipelines

- Suite of plug-ins to support Continuous Delivery pipelines
 - Jenkinsfile (in Git): parameters, stages, steps
 - Declarative (more recent, powerful) and scripted syntax
- Why?
 - To control docker/non-docker builds (linux docker, arm, mac) in same job
 - To get rid of many independent Jenkins jobs
 - To have Jenkins configuration in Git
- Both full nightlies and releases delivery have been re-implemented in pipelines, in two jobs:
 - lcg_release_pipeline: [jenkinsfile-release](#)
 - lcg_nightly_pipeline: [jenkinsfile-nightly](#)
- Single matrix job for all [nightlies](#) and for all [releases](#)
- Also for specific projects:
 - VecGeom , AdePT (in progress)

Jenkins pipelines screenshot

Stage View		Environment	InDocker	Build	CopyToEOS	CreateRPMs	InBareMetal	Build	CopyToEOS	MacPreinstall	CVMFSInstall	LCGinfo	PublishRPMs	TestRPMs
Average stage times: (Average <u>full</u> run time: ~11min)		1min 32s	18s	1h 26min	2s	35ms	1min 3s	0ms	0ms	58s	2min 33s	31ms	30ms	31ms
#122 geant4ext20210118-mac1015-nativ Jan 19 15:31 No Changes		1min 5s					4min 12s	5min 51s	1s	3min 52s	5min 33s			
#121 geant4ext20210118-ubuntu20-nativ Jan 19 15:31 No Changes		1min 4s	7s	1min 29s	2s						4min 2s			
#117 geant4ext20210118-centos7-gcc8 Jan 18 22:20 No Changes		39s	5s	4min 14s	6s						2min 21s			

Last year activities: highlights (2)

- **LCGCMake**
 - Commissioning of the *Layered Stack* concept
 - In production since LCG_97
 - Build/Deployed layers for ATLAS, LHCb and FCC
 - Improved caching of binaries through *latest*
 - Consolidated usage of binaries artifacts
 - Speed-up of nightlies
- **Python 3 now the default**
 - Starting from LCG_99
 - As of 11/2020, nightlies dev{3,4} are python3
 - ATLAS layers are Python3 only
 - Python 2 builds kept to help transition
 - dev{3,4}python2

Layered Stacks

- Address issues with multiple versions of packages in same build
 - Also with adding new version of packages (e.g. MC generators)
- Layered stacks:
 - A stack/layer contains a **single version** of each package
 - Layers can be **build incrementally on top of each other**
 - Allowing removal / change-of-{version, recipe} of packages
 - Allows to **tailor** a give release to **specific needs**
- Current status, documented in [lginfo](#)
 - LHCb: pro version LCG_97a
 - Layers LCG_97a_LHCb_{1,2,3,4}; tuned MC versions
 - ATLAS: pro version LCG_98python3
 - Layers LCG_98python3_ATLAS_{1,2,3,4,5}; build only required packages and deps (~150)
 - FCC: pro version LCG_97a
 - Layers LCG_97a_FCC_{1,2,3,4}; include FCC specific packages

Last year activities: highlights (3)

- **Testing**

- Gaudi tests run in dev3 for Gaudi HEAD against ROOT HEAD
 - Attempt to anticipate problems as the one observed during migration to new PyROOT
- Prototyped running of roottest

- **Package content**

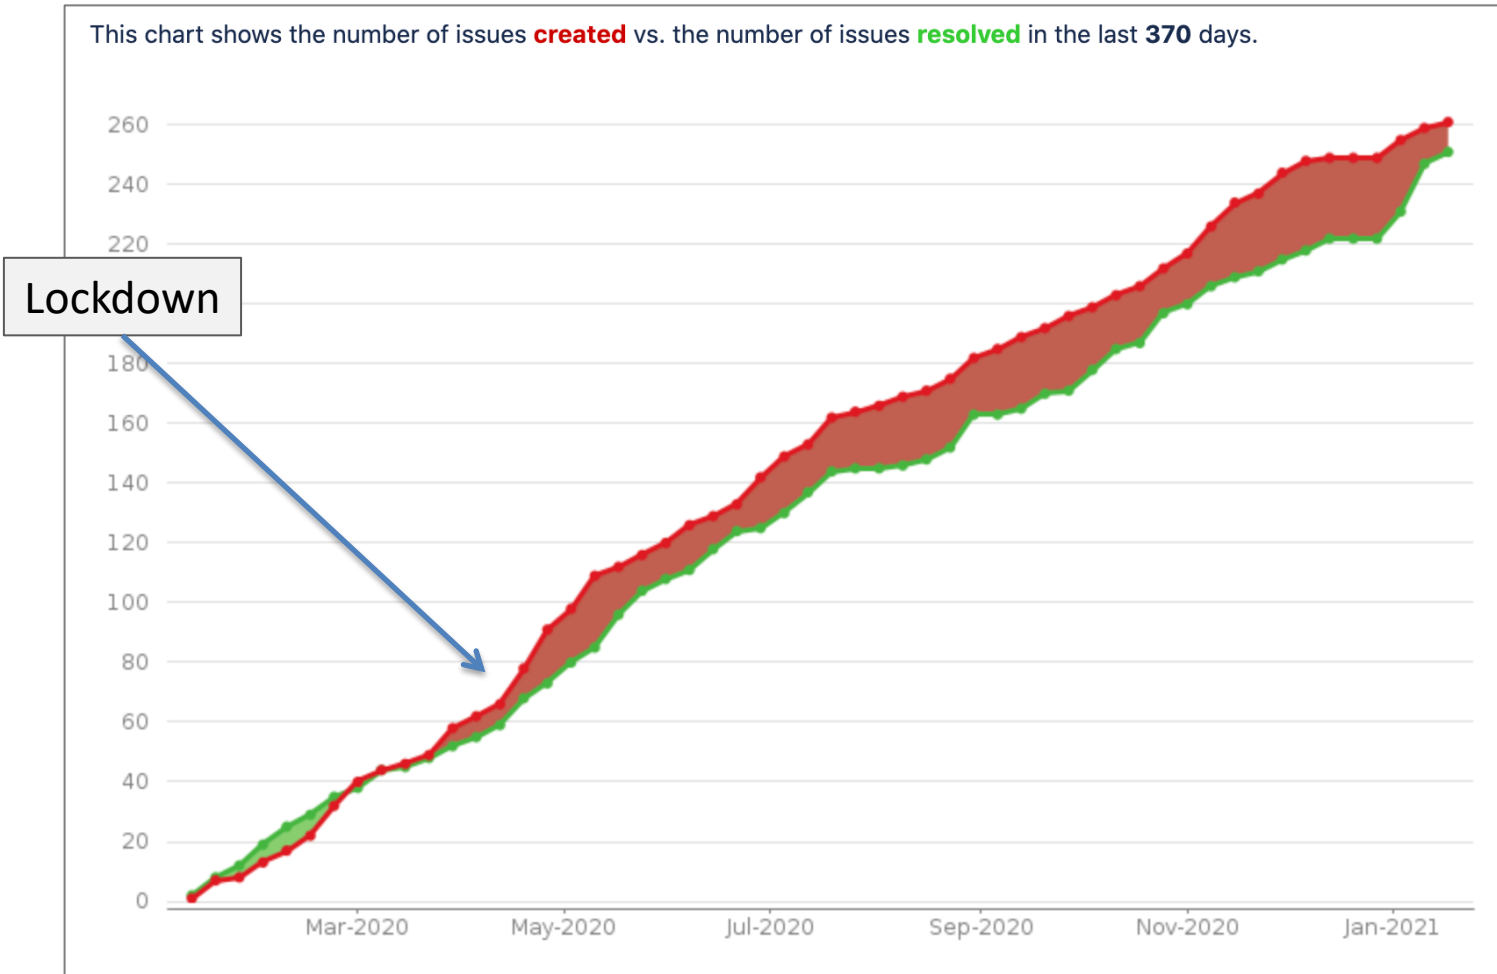
- Continued to critically review the package content / versions
 - Considerable fraction of work goes into package updates
- New packages of general interest
 - Gaudi, DD4hep, EDM4hep, acts, Podio, ... (for key4hep)
- Many new generators / versions included

LCG_96b → LCG_99
+ 81 new packages
- 12 removed
134 updated packages

- **Support**

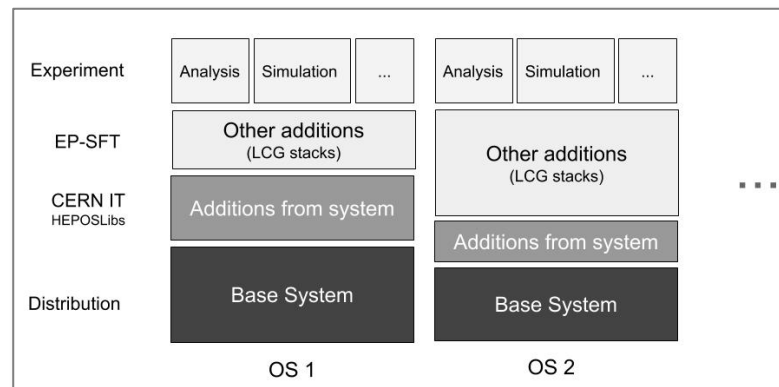
- [JIRA](#) continued to be the **main support channel and planning tool**
- Service Now Functional Element: Software Development For Experiments
 - Low activity (1-2 /month) mostly connected to AFS phase out

JIRA created/resolved tickets



Last year activities: highlights (4)

- [HepOS-Libs](#) or the problem of *defining the baseline system*
- Dedicated LIM on 12 May 2020



*There a feeling that **HEP_OSlibs won't be needed in the long term**. However, **the situation is not mature** for that. For the time being (CentOS 8 and Ubuntu 20.04) we proceed as before, with meta-packages provided as usual. The definition of these will wait the first successful LCG builds to make sure they contain all new package which may be needed.*

Migrating away from Drupal 7

- Reminder: CERN revised D7 phaseout strategy
 - D8 recommended, with in-house support (FELL/TECH/IT) or outsourced
 - Opening to alternative, simpler, solutions (e.g. Jekyll)
 - Migration deadline (shutdown of D7): 31 August 2020
- Dedicated resources (A Stano, fellow) for EP *connected* sites
 - [EP-DEP](#), [EP-DEP-AGS](#), [EP-DEP-ESE](#), [EP-DEP-DT](#), [EP-DEP-SFT](#), [EP-RND](#), [EP-NEWS](#), [FCC-EE](#), [IML](#), [NEUTRINO-SECRETARIAT](#), [USEROFFICE](#), [ACCU](#), [NEWCOMERSGUIDE](#), [AGS-ADMIN-GUIDE](#), [GEANT4](#)
- Some projects decided to move away from Drupal (Jekyll)
 - [CernVM](#), [CernVM-FS](#), [SPI](#) (new site), [ROOT](#)

Last year activities: highlights (5)

- Spack

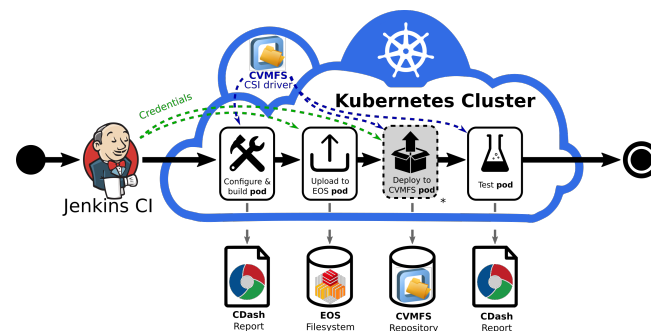
- Not much progress for the LGC stack itself
 - Summer Student project cancelled
 - Required Spack developments were delayed (new concretizer)
 - Lack of immediate need by experiments
- Advise/help provided to key4hep Spack instance

- Kubernetes

- Investigations of Kubernetes continued
 - Complete prototype based on ARGO workflow engine
 - Missing piece CernVM-FS pod
- R Bachmann technical studentship

- AoB

- Grafana: IT provided MONIT service sufficient for basics needs
 - Phase-out EP-SFT Grafana instance / service



Last year activities: releases

Releases

- LCG_97{a}, python3: ROOT v6.20/02{06}
 - Base layers for FCC, LHCb pro release (LCG_97a)
- LCG_98, python3: ROOT v6.22/00
 - Base layers for ATLAS pro release
- LCG_99: ROOT v6.22/06
 - [default python3](#)
 - Ports on CentOS8, Ubuntu 20 LTS, MacOSX; drop SLC6

Special 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, used in SWAN
- devARM
 - Port on ARM, used by LHCb and ATLAS

Plan Of Work for 2021

- **Infrastructure**
 - Consolidation, investigate Singularity, GitLab-CI
- **Build system**
 - Keep stack content up-to-date
 - LCGCMake streamlining
 - Spack concrete progress
- **Improvements / Developments**
 - Kubernetes
- **Support**
 - Users, new tools / platforms, ...
- **Documentation**

Infrastructure

- Jenkins
 - Consolidation of pipelines builds
 - Maximize use of parallelization, optimize scripts
 - Clean-up of obsolete jobs
- Investigations
 - Integration with k8s
 - Pipelines functionality
 - GitLab-CI
 - Singularity for build nodes
 - Simplified build node setup / management
 - Better integration with CernVM-FS served repositories
 - `unpacked.cern.ch`

Improvements: LCGCMake

- Streamline usage of single source for binaries
 - Use tarballs for CernVM-FS and RPMs
 - Automate replacement of single package in release
- Improve tagging of builds
 - Prototype /commission the introduction of *instruction set* family name in tags
- Re-discuss number of digits in compiler tag
 - Several incompatible effects due to MINOR changes
 - e.g. gcc8.2 vs gcc8.3
 - To be discussed in LIM

Improvements: RPM, testing

- RPMs
 - Consolidate 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 the possibility to get rid of meta-RPMs
 - Replace with scripts using info from .txt files, as done by LHCb
- Testing
 - Increase coverage for releases or nightlies
 - Finalize inclusion of *roottest* in regular tests
 - Specific pipeline step
 - Run other similar test suites?
 - Integration tests provided externally (e.g. experiments)
 - GaudiTest as example

Improvements: CernVM-FS publication

- Minimize the impact of this step by **maximize concurrency**
 - Speed up standard publication step
 - CernVM-FS publication POD in k8s setup
- **CernVM-FS 2.8.0** contains the relevant ingredients
 - `cvmfs_server` enter writing shell
 - CernVM-FS Conveyor high-level interface
- Plan is to develop/test using the **HSF oriented** repositories
 - `sw.hsf.org`, `sw-nightlies.hsf.org`
- In collaboration with CernVM core team

See also [CernVM-FS high-level publication interface](#)

Developments

- Spack
 - Concretise investigation into a **usable test build** taking advantage of
 - Latest Spack developments
 - Experience acquired with key4hep
 - Coordinate with HSF Software Tools and Packaging WG
- Containers
 - Continue to investigate Kubernetes-based build service
 - Investigate Jenkins pipelines k8s integration
 - AK technical student project
 - Streamline production of release containers for tutorials, ...
 - Investigate use of CernVM-FS shrinkwrap to optimise images
 - Provision of *baseline-system-definition* containers
 - Ready-to-use minimal image with CernVM-FS, EOS

Support / Documentation

- Support
 - [LCG stack customers](#)
 - And users of /cvmfs/sft.cern.ch (as AFS replacement)
 - Follow requests for [new architecture, platforms, compilers, tools](#)
 - PowerPC, ARM, ...
 - In coordination with IT
 - Follow requests for [new packages and versions](#)
- Documentation
 - [lcginfo](#): front-end for release information
 - Consolidate browsing functionality
 - Upgrade Django framework
 - [lcgdocs](#): SPI documentation
 - Streamline content organization

AoB

- Follow situation with CentOS8
 - Adapt to new CERN IT model under discussion
- Develop strategy to attract new clients
 - ITUM presentation
- Workshop with stakeholders and users to get feedback
 - Or dedicated extended LIM

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	

Please volunteer!

Service Tasks remarks

- *Grafana* phase-out
- *Other OS nodes* only used by ROOT
- *Printer Room HW*
 - New entry for minimal maintenance operations ?
 - Powercuts, etc etc

Outlook for 2021

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