CernVM, Systems & Services Plan Of Work for 2019

G Ganis, J Blomer

On behalf of the CSS team

21 January 2019

Outline

- Quick reminder about the CSS team
  - Team members
- System & Services
  - Last year activities
  - A plan of work for 2019
- CernVM
  - J Blomer

PoW 2018 presentation

- G Ganis, CernVM, System & Services Plans for 2018, 12 Mar 2018

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 software distributed
- Preserve and develop activities and/or products rotating around it

CSS Manpower situation

		2019	2020	
G Ganis	St	80%	80%	C, S
J Blomer	St	50%	50%	C
I Goulas	St	100%	100%	S
E Obreshkov	St	10%	10%	S
S M Muzaffar	St	10%	10%	S
R Popescu	Fl	66%		C, S
A Stano	Fl	80%	20%	W
J Heinz	TS	66%		S
S Mosciatti	TS	25%		C, S
J Priessnitz	TS	60%	60%	C
I Razumov	PAS	40%		G, S
Russian WLCG	User	100%	100%	G, S
Total (FTE) (7.95 in 2018)		6.85+	4.3+	

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

ATLAS and CMS librarians hosted by SFT

Valuable contribution from P Mato

Team members activity reports in 2018

CernVM

- S Mosciatti, Ingesting tarballs to integrate containers in CernVM-FS, 11 Jun 2018
- R Popescu, A High Level publication interface for CernVM-FS, 3 Dec 2018
- R Popescu, CernVM Users Workshop Report, 19 Feb 2018

SPI

- R Pacholeck, Technical Studentship Overview, 27 Aug 2018
- P M Lorenzo, Summary and Conclusions from the LIM workshop, 2 Jul 2018

Web

- N Kousi, The SFT web site and Drupal 8: the state of things, 14 May 2018

AoB

- R Popescu, Programming languages for frameworks: is C++ still the best choice?, 23 April 2018
- J Blomer, ITUM Summaries: 12 Nov 2018, 4 June 2018, 19 Feb 2018

System & Services
(aka SPI)

S&S deliverables and stakeholders

Deliverables:

- Software stack of about 400 packages
  - {slc6, cc7, ub16, ub18, …}x{gcc62, gcc7, gcc8, clang, …}x{opt, dbg}
- Releases under /cvmfs/sft.cern.ch (C) and as RPMs (R) and tarballs (T)
- Nightlies under/cvmfs/sft.cern.ch

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

S&S stakeholders

- ATLAS, LHCb: nightlies (C), releases (R)
- FCC, SWAN: nightlies and releases (C)
- BE: releases (C); possibly R or T
- NA62: releases (C)

- Many users on lxplus and elsewhere

Last year activities

Main items:

- Consolidation and maintenance of the infrastructure
- Improving testing
- R&D contribution to Spack
- L&I workshop to collect input from stakeholders

L & I workshop

- 30 May 2018, 31 registered people
  - All main stakeholders represented
- Good dissection of the current situation
- Spot problems and areas of improvement
  - Testing
  - RPMs
    - Need for release candidates
  - Strategy for package content (addition / retirement)
  - Performance of the whole build and release process
    - Impact of chosen hardware / solutions
  - Support
    - Need for a clear support channel
    - Essential documentation
  - Facilitate external contributions
    - Merge request testing and validation

L & I workshop (cnt'd)

- Role / future of HepOSLibs
  - Definition of base platform
- Platform support
  - Improved ARM, possibly Power
  - MacOsX support required by several customers
- Role of containers

Last year activities: releases

Releases

- LCG_93{a,b,c}, LCG_93python3: ROOT v6.12
- LCG_94{a}, LCG_94python3: ROOT v6.14

New views

- LCG_94python3_nxcals
  - Dedicated view for NXCALS usage in SWAN for BE
- dev3cuda9
  - Build with CUDA support for TensorFlow, …

Last year activities: highlights

- Build of *compilers* and *binutils* integrated in LCGCmake
- *Release Candidate* concept exercised with 94 (LCG_94rc1)
  - Quite positive; needs to improve coverage of testing by the experiments
- Consolidation of dev toolchains (dev-base)
- Fixed Python dependencies
- Improved nightly-report tool
  - Based on Jenkins information
- Improved consistency of lcginfo information
- Simplified provision of new container images
- Support for micro-architectures in LCGCmake
- Quick start mode for LCGCmake (bin/lcgcmake command)
- Many new generators / versions included

Last year activities: highlights (2)

- Support
  - JIRA re-established as main support channel
  - Weekly shifts to monitor status of builds and reduce the reaction time to problems
  - Re-connected to dedicated *Service Now* functional element
    - *Software Development for Experiments*
- R&D contribution to Spack (HSF context)
  - Summer Student Project:
    Optimizing Experiments' Software Stack Management With Spack
    - P Chelarescu (supervisors: J C Villanueva, P M Lorenzo, G Stewart)
    - Reports at SFT Group Meeting, HSF Packaging Group
  - Positive evaluation of *Spack ccache* and *Spack chain*
    - Significant speed-up of builds

From last year PoW

SPI

Consolidation and maintenance

- LCG build and deployment infrastructure [3m] ☑
  - Tools and scripts optimizations
  - New platforms and compilers as requested
    - AVX2, ARM64
- Consolidate release for the Beams department [1m] ≈
  - LCG…_python3
- Review the list of packages [1m] ☒
  - Automation of package versions and builds
    - GitHub notifications

GenSer

Maintenance

- Continue fulfilling requests from the experiments and users for addition of new generators and/or new versions ☑

Developments ☒

- Migration to LCGTest
- Investigate the possibility to use T4T infrastructure for generator testing ≈

SPI

Consolidation and maintenance (cont'd)

- Sort the /cvmfs/sft.cern.ch/ top area [2m] ≈
  - Remove unused (or misplaced) directories, AFS references, …
  - Document the content and the way to use it
    - Add README, promote views, lcgenv

Development, Commissioning

- Bring LCGTest in production state [3m] ☒

R&D

- Incremental builds w/ Spack [3m] ☑
  - Summer Student

October's brainstorming about SPI re-organization

- 1 day discussion (notes)
- Main items
  - Re-define main deliverable and services
  - Streamline and Optimize Procedures
  - Support Process
  - Re-define strategies to minimise number of packages and platforms
  - Redistribution of roles
  - Longer term strategy and connection with HSF

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

Infrastructure

- Jenkins
  - Server upgrade to version 2.160.3 (now 2.46.3) [Q1]
    - Bug fixes and (many) security fixes
  - Major upgrade: some plug-ins need different settings
    - Testing on a separate, puppetized, instance
  - Nodes [Q1]
    - Evaluate alternative staging area (e.g. NFS, CephFS)
    - Move away HOME from AFS
- Consolidate use of containers
  - Debugging of failed jobs [Q1/Q2]
  - Investigate use Kubernetes for an optimized build service [Q4]
    - JH master thesis

Improving CernVM-FS publication

- Use of CernVM-FS high-level publication interface
  - Repository Gateway, RabbitMQ, Job DB
  - Allows to optimize operations, handling dependencies and maximizing concurrency
  - Inspired from LHCb's nightly build publishing system
- Requires new repositories, CernVM-FS Gateway-enabled
  - sw.hsf.org, sw-nightlies.hsf.org
- Start with parallel publishing

[Q1/Q2]

Improving RPM repository

- Targets:
  - Standardize repository
  - Reduce size of a single database

[Q1/Q2]

- Use OS, ARCH, versioning
  - Currently:
    `ROOT-a1638_6.14.08_x86_64_centos7_gcc8_dbg-1.0.0-94.noarch.rpm`
- Investigate debuginfo technology

Improve testing

- Validate as much as possible a release or a nightly
- LCGTest: currently enabled to test import of python modules and their dependencies
  - Difficult to extend to other tests
- Target
  - Run basic functionality (e.g. --version) and consistency checks [Q1/Q2]
  - Run more complex test (e.g. roottest) on installed package
  - Open possibility to run integration tests provided externally
    - E.g. from experiments

Review package content

- Strategy to deprecate packs
  - Keep a package until it gives troubles
    - Then find out if still needed
  - Review content and versions at each new major release [Q1-Q4]
- Debug builds
  - More than duplicate the need for resources
    - Really needed only for subset of packages: ROOT, XRootD, …
  - Determine list of packages for which symbols are needed
  - Setup smooth use of mixed builds [Q2]

Web sites move to Drupal 8

- In 2018 migration to Drupal8 suffered from blocking instabilities in the testing infrastructure
  - Priority was given to the main CERN site
  - Kit for site migration only available in December
- Current plan of work
  - Move sites SFT (w/ some revision), ROOT, other SFT (CernVM, G4), DT, EP, EP-AGS
  - JIRA SPI-1010 to SPI-1019
- Fellow allocated (first year on EP budget)

[Q4]

New ports, Documentation

- ARM
  - Interest from several stakeholders
    - Starting point: initial port from LHCb
    - Possible Summer Student project
  - Need to secure hardware
    - TechLab? Own Servers?
- Documentation
  - LCGinfo
    - Add precise information per platform [Q1]
  - lcgdocs
    - Document all relevant component and actions [Q4]

AoB

- Dedicated release-like view for BE
  - Limited package content (~100 packages)
  - Detached from main LCG release (no ROOT)
    - E.g. LCG_BE_1
- AFS phase-out of /afs/cern.ch/sw/lcg
  - Discussed at LIM on 11.12.2018
    - 2019-02-01: make inaccessible (except for known usage: mostly ATLAS)
    - 2019-06-01: remove any still-inaccessible areas
    - 2019-12-01: last negotiated end dates for temporary access on AFS
- Follow HSF Packaging Forum activities

Service Tasks

Service/Task	Main Responsible	Alternate	Documentation
Jenkins service	Shahzad Muzaffar	Gunter Folger	HowTo
Coverity service	(Gabriele Cosmo)	Axel Naumann	
CDash service	Javier Cervantes	Emil Obreshkov	HowTo
SLC/CC nodes	Shahzad Muzaffar	Gunter Folger	
Windows nodes	Bertrand Bellenot	Gunter Folger	
Mac nodes	Radu Popescu	Witold Pokorski	
Other OS nodes	Gerardo Ganis	Axel Naumann	HowTo
Drupal Manager	Andrea Stano	Danilo Piparo	
ITUM contact	Jakob Blomer	Graeme Stewart	
C5 contact	Jakob Blomer	Enric Tejedor	HowTo
Jira Service	Ilias Goulas	Emil Obreshkov	
Training	Danilo Piparo	Enric Tejedor	

Need an effort to provide documentation!
Perhaps as markdown files in GitLab

Summary

- For S&S, 2018 has been a year of restructuring and discussion
  - Redistribution of roles due to manpower changes
  - Identified area of improvements and consolidation
    - Also exploiting in house expertise and products
- Towards a sustainable, for the group, Software Provisioning Service