

# UPDATE OF THE LINEAR COLLIDER

# Software Infrastructure

A Love Story between a Parrot and a Whale

Marko Petrič



CLIC Workshop 2017

*CERN, 7 March 2017*

# Outline

- ▶ AFS phaseout
- ▶ The Past
- ▶ Current status of Infrastructure
- ▶ Plans and open issues
- ▶ New hardware resources

# Migration AFS → CVMFS and EOS

- ▶ Concern for the long term viability of the OpenAFS project
- ▶ Maintenance continues but little new investment
- ▶ Phase out AFS by LS2

AFS	CVMFS	EOS
Webpages		✓
iLCSoft releases	✓	
iLCSoft tar balls		✓
Test beam data		✓
DIRAC	✗	
Generic software	✓	

- ▶ DIRAC: establish everyone can access `/cvmfs/clicdp.cern.ch/`
- ▶ Some legacy activity still in AFS, but no write activity

# The Past (1/2)

## ABOUT

You can customize this short message in the `index.tmpl` of this template in order to tell your visitors what they find in your repositories.

Visit [www.websvn.info](http://www.websvn.info) for more information about WebSVN.

Learn more about Subversion at [subversion.tigris.org](http://subversion.tigris.org).

## SUBVERSION REPOSITORIES

▼ EuXFEL		
▼ Editing	Rev 0	
▼ mtca4u_applications		
▼ ad16	Rev 79	304d 09h
▼ DOOCS_Adapter	Rev 145	284d 04h
▼ matlab_tools	Rev 113	289d 07h
▼ General		
▼ afsadmin	Rev 337	275d 06h
▼ afsbackup	Rev 31	372d 03h
▼ aidasoft	Rev 2338	166d 06h
▼ alfa_sw	Rev 27	627d 08h
▼ alibavasoft	Rev 221	583d 00h
▼ astrofit	Rev 207	1018d 03h
▼ AtlasSiliconUpgrade	Rev 184	221d 07h
▼ b2vxd_testbeam	Rev 238	3d 07h
▼ Baskaran	Rev 10	1277d 06h
▼ bbq	Rev 135	99d 08h
▼ bbrtrack	Rev 8	3825d 05h
▼ BCM1F	Rev 231	821d 02h
▼ benchmark_yangsuli	Rev 86	3091d 11h
▼ boojum	Rev 8	1151d 08h
▼ brahms	Rev 91	4062d 22h



- ▶ All software hosted at DESY
- ▶ Distributed via SVN
- ▶ Was good solution for its time
- ▶ Lacking several modern attributes:
  - ▶ Uncontrolled committing
  - ▶ No review of code
  - ▶ No testing
  - ▶ No handle on regressions
  - ▶ No diagnostics

# The Past (2/2)

- ▶ Testing via nightly builds
- ▶ Only detect errors after buggy code has been injected



## CDash projects

### Dashboards

Project	Description	Last activity
DD4hep	Detector Geometry Description for HEP	13 hours ago
DD4hepExamples	Examples shipped with DD4hep - built as a standalone project for testing.	11 hours ago
lcgeo	Lineaer Collider detector geometry description based on DD4hep	12 hours ago

[Show all 23 projects](#)



CDash 2.0.2 © Kitware | [Report problems](#) | 1.5Gb

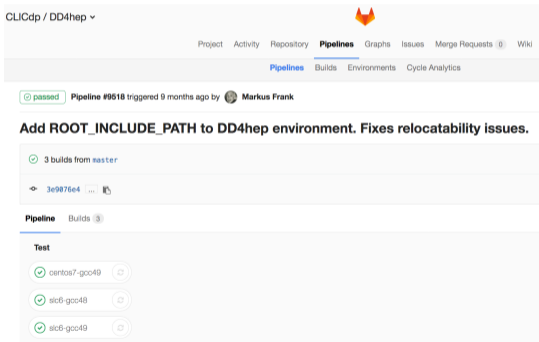
### Nightly

Site	Build Name	Update	Configure		Build		Test			Build Time
		Files	Error	Warn	Error	Warn	Not Run	Fail	Pass	
ilcsoft_build_sl6.desy.de	Linux-2.6.32-573.12.1.el6.x86_64-/afs/cern.ch/sw/lcg/contrib/gcc/4.9.3/x86_64-slc6/bin/g++-Debug		0	0	0	50	0	0	10	14 hours ago

# The DD4hep Showcase

## Immediate testing is good

- ▶ New computing resources dedicated to testing
  - ▶ 12 × 4 code instances, 2 GB / core RAM
- ▶ Mirror DD4hep every few minutes to CERN/Gitlab and run tests



CLICdp / DD4hep

Project Activity Repository **Pipelines** Graphs Issues Merge Requests Wiki

Pipelines Builds Environments Cycle Analytics

passed Pipeline #9618 triggered 9 months ago by Markus Frank

**Add ROOT\_INCLUDE\_PATH to DD4hep environment. Fixes relocatability issues.**

3 builds from saster

3e9876e4

Pipeline Builds 3

Test

- centos7-gcc49
- slc6-gcc48
- slc6-gcc49


```
--- Removed runtime path from "/builds/CLICdp/DD4hep/bin/dumpdetector"
--- Installing: /builds/CLICdp/DD4hep/bin/test_surfaces
--- Removed runtime path from "/builds/CLICdp/DD4hep/bin/test_surfaces"
--- Installing: /builds/CLICdp/DD4hep/lib/6__teve1_rdict.pcm
--- Installing: /builds/CLICdp/DD4hep/lib/6__teve1_rdict.pcm
--- Installing: /builds/CLICdp/DD4hep/bin/teveLCIO
--- Removed runtime path from "/builds/CLICdp/DD4hep/bin/teveLCIO"
--- Installing: /builds/CLICdp/DD4hep/bin/teveDisplay
--- Removed runtime path from "/builds/CLICdp/DD4hep/bin/teveDisplay"
$ .. /bin/thisdd4hep.sh
$ make test
Running tests...
Test project /builds/CLICdp/DD4hep/build
Start 1: t_test_example
1/9 Test #1: t_test_example ..... Passed 0.88 sec
Start 2: t_test_units
2/9 Test #2: t_test_units ..... Passed 2.17 sec
Start 3: t_test_surface
3/9 Test #3: t_test_surface ..... Passed 2.00 sec
Start 4: t_test_bitfield64
4/9 Test #4: t_test_bitfield64 ..... Passed 0.88 sec
Start 5: t_test_DetType
5/9 Test #5: t_test_DetType ..... Passed 0.88 sec
Start 6: t_test_PolarGridRPhI2
6/9 Test #6: t_test_PolarGridRPhI2 ..... Passed 0.89 sec
Start 7: t_test_cellDimensions
7/9 Test #7: t_test_cellDimensions ..... Passed 0.90 sec
Start 8: t_test_cellDimensionsRPhI2
8/9 Test #8: t_test_cellDimensionsRPhI2 ..... Passed 0.89 sec
Start 9: t_test_EventReaders
9/9 Test #9: t_test_EventReaders ..... Passed 2.18 sec
```

- ▶ Introduced “real time” testing on SLC6 and CC7 with gcc 4.8(9)

# The DD4hep Migration (1/2)

Testing before is even better

- ▶ Move DD4hep repository from DESY to Github
- ▶ Use Pull Requests to put code into repository

 **Add deprecated attributes** ✓

#117 by andresailer was merged 2 days ago

 2


 **Fix problem with DetElement transformations.** ✓

#115 by MarkusFrankATcernch was merged 3 days ago


 12

 **Add GuineaPigReader from lcggeo** ✓

#113 by andresailer was merged 4 days ago

 **cmake: set\_compiler\_flags: improve checking of variables** ✓

#112 by andresailer was merged 4 days ago

 **Update to allow the externalization of the XML and the ComponentPrope...** ✓

#111 by MarkusFrankATcernch was merged 6 days ago

 **multiple vertices per interaction continued** ✓

#110 by gaede was merged 5 days ago

 4

# The DD4hep Migration (2/2)

## Testing before is even better

- ▶ Code review
- ▶ Utilise Travis-CI to run test
  - ▶ 2 Cores and 40 min available for free

```
514 -- Configuring done
515 -- Generating done
516 -- Build files have been written to: /DD4hep/build
517 [155/521] Building CXX object DDCore/C...ore.dir/src/VolumeManagerInterna.cpp.o
518 ../DDCore/src/VolumeManagerInterna.cpp: In constructor
519 'DD4hep::Geometry::VolumeManagerContext::VolumeManagerContext()':
519 ../DDCore/src/VolumeManagerInterna.cpp:26:32: warning: 'DD4hep::Geometry::VolumeManagerContext::mask' is
depreciated: This member variable might get axed if it is not used, please tell us if you do [-Wdepreciated-
declarations]
520     : identifier(0), mask(~0x0ULL) {
521         ^
522 In file included from ../DDCore/src/VolumeManagerInterna.cpp:17:0:
523 ../DDCore/include/DD4hep/objects/VolumeManagerInterna.h:61:16: note: declared here
524     VolumeID mask;
525         ^~~~~~
526 ../DDCore/src/VolumeManagerInterna.cpp:26:32: warning: 'DD4hep::Geometry::VolumeManagerContext::mask' is
depreciated: This member variable might get axed if it is not used, please tell us if you do [-Wdepreciated-
declarations]
527     : identifier(0), mask(~0x0ULL) {
528         ^
529 In file included from ../DDCore/src/VolumeManagerInterna.cpp:17:0:
530 ../DDCore/include/DD4hep/objects/VolumeManagerInterna.h:61:16: note: declared here
531     VolumeID mask;
532         ^~~~~~
```



gaede commented 18 days ago

Member

@MarkusFrankATcernch please check logic in [c6a818f](#).



petricm reviewed 18 days ago

View changes

DDG4/lcio/LCIOEventReader.cpp

Hide outdated

```
... .. @@ -146,11 +146,10 @@ LCIOEventReader::readParticles(int event_number,
146 146     vtx->y = p->vsy ;
147 147     vtx->z = p->vsz ;
148 148     vtx->time = p->time ;
149 -     haveVertex = false ;
149 +     haveVertex = true ;
```



petricm 18 days ago

Member

This does not change anything. In line 86 `bool haveVertex = true`; subsequently if the `if( !haveVertex && genStatus == 1 ){` is always false as describes in CID 1402993. You have to change `if( !haveVertex && genStatus == 1 ){` to `if( haveVertex && genStatus == 1 ){` to address this.

- ▶ Due to time constrains from Travis run only type of build SLC6 with gcc 6.2

# Implementation

- ▶ Travis-CI offers Ubuntu, but we don't want this
- ▶ Use Dockerized version of CernVM
  - ▶ Docker → additional layer of operating-system-level virtualization
  - ▶ CernVM + parrot connector to access CVMFS
  - ▶ All dependent software in CVMFS (quick migration)
- ▶ Due to time limit only try to build software in Travis-CI
- ▶ Post merging sync to Gitlab and run more tests
  - ▶ No time limit, use own 12 instances
  - ▶ Again use Docker for different systems SLC6 and CC7
  - ▶ Test different compilers GCC 6.2 and LLVM/Clang 3.9
  - ▶ Purchased Mac hardware to test on Mac
  - ▶ All dependencies centrally in CVMFS
- ▶ Successfully established continuous QA for DD4hep → iLCSoft

# iLCSoft migration

iLCSoft	Build Status	Oxygen
CEB	Build unknown	
CEBViewer	Build unknown	
CLICPerformance	Build pending	
Clusetta	Build pending	
CondDBMySQL	Build unknown	
ConformalTracking	Build unknown	
DDKaTar	Build unknown	
DDMariePandora	Build pending	
EasyJetClustering	Build unknown	
ForwardTracking	Build unknown	
GEAR	Build unknown	
iLCInstall	Build unknown	
iLCList	Build unknown	
iLCConfig	Build unknown	
iLCPerformance	Build unknown	
KaITest	Build unknown	
KaITest	Build unknown	
KITrack	Build unknown	
KITrackMatrix	Build unknown	
LCCD	Build unknown	
LCTPlus	Build pending	
LCTVertex	Build unknown	
logoo	Build pending	
LCD	Build pending	
LCTuple	Build unknown	
MacIn	Build pending	
MacInDD4hep	Build pending	
MacInFastJet	Build pending	
MacInKfkt	Build unknown	
MacInKfktProcessors	Build unknown	
Overlay	Build pending	
FCALSW	Build Status	Oxygen
FCALClusterer	Build pending	

- ▶ After success of DD4hep do same to iLCSoft
- ▶ Identify all code authors for whole iLCSoft history
- ▶ Use GPL v3 to license code
- ▶ Packagewise migrate to <https://github.com/iLCSoft/>
- ▶ Migrated 13 of 34 packages
- ▶ To test package need “recent” version of iLCSoft due to dependencies
  - ▶ Nightly build of iLCSoft deployed to CVMFS
  - ▶ Use same Travis-CI mechanism as in DD4hep
- ▶ Elimination of all warnings from several packages
- ▶ Ensure better code quality
- ▶ Check status of migration at <http://ilcsoft.github.io/>

# iLCSoft builds

- ▶ For rapid development enable snapshot build
- ▶ Not real releases, but status on date
- ▶ Available under `/cvmfs/clicdp.cern.ch/iLCSoft/builds/`

2016-02-02	2016-07-04	2016-09-27	2016-11-24	2017-02-17
2016-04-06	2016-08-22	2016-11-09	2016-12-15	current
2016-06-22	2016-09-12	2016-11-22	2017-02-02	nightly



- ▶ All builds available in iLCDirac for mass production
- ▶ Access to nightly build for development

# Code Analysis

- ▶ Open Source → free access to Coverity
- ▶ Coverity is a static code analysis and dynamic code analysis tool
- ▶ Pair all packages with Coverity and examine defects

```
4. overflow: Add operation overflows on operands 9UL and siz_nam. Example value for operand: siz_nam = 18446744073709551607.  
◆ CID 1398170 (#3-1 of 4): Overflowed array index write (INTEGER_OVERFLOW)  
5. overflow_sink: Overflowed or truncated value (or a value computed from an overflowed or truncated value) 9UL + siz_nam used as array index.  
159 text[8] = text[9+ siz_nam] = text[10+ siz_nam+ siz_add] = 0;
```

- ▶ DD4hep: eliminated 138/171 code defects




- ▶ Due to new structure harder to introduce new defects
- ▶ Already deployed to some packages

# Future Plans

- ▶ Transition all code to Github
- ▶ Equip all packages with CI and code analysis
- ▶ Integration builds testing on commit basis
- ▶ Memory consumption analytics
- ▶ Extension of the test suite

# New Hardware resources

- ▶ Investment in Deep Learning system
- ▶ 2x Intel Xeon (40 Cores), 4x GeForce GTX 1080, 256GB RAM and 2 TB disk



- ▶ System in deployment, will be made accessible to collaboration members with appropriate projects (calorimetry, flavour tagging...)
- ▶ Procedures to be communicated once the system is in production

# Summary

- ▶ Transitioned nearly all software to CVMFS (AFS deprecated)
- ▶ Transitioned DD4hep and iLCSoft to Github
- ▶ Imposed comprehensive testing system
  - ▶ Code review
  - ▶ Testing
- ▶ Improved code quality and minimized regressions
- ▶ Work to migrate all iLCSoft still ongoing
  - ▶ Scheduled to be finished in LC Hands On Week
- ▶ Testing infrastructure open to all members (used by allpix2)
  - ▶ Contact us if you want to take part
- ▶ New GPU hardware soon to be available to users for physics analysis