

Installation of CMSSW in the Grid

DESY Computing Seminar

May 17th, 2010

Wolf Behrenhoff, Christoph Wissing



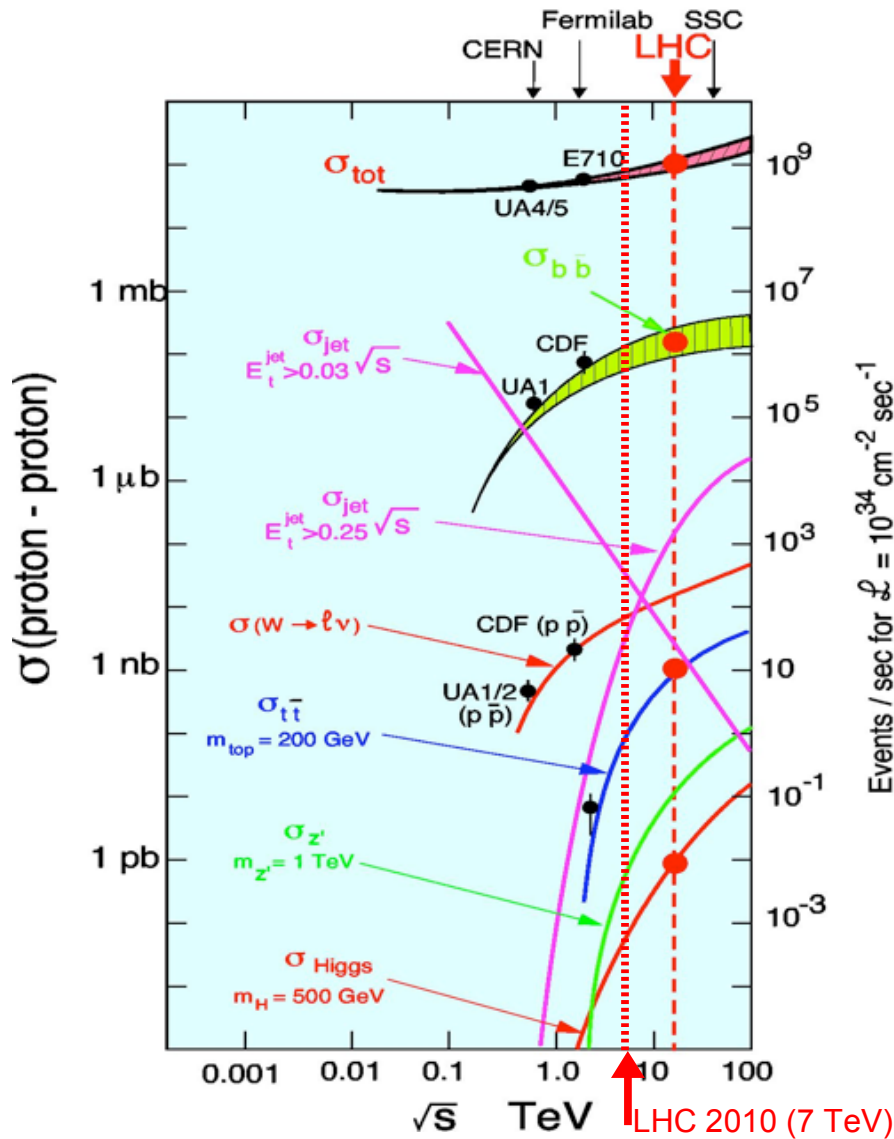


Contents

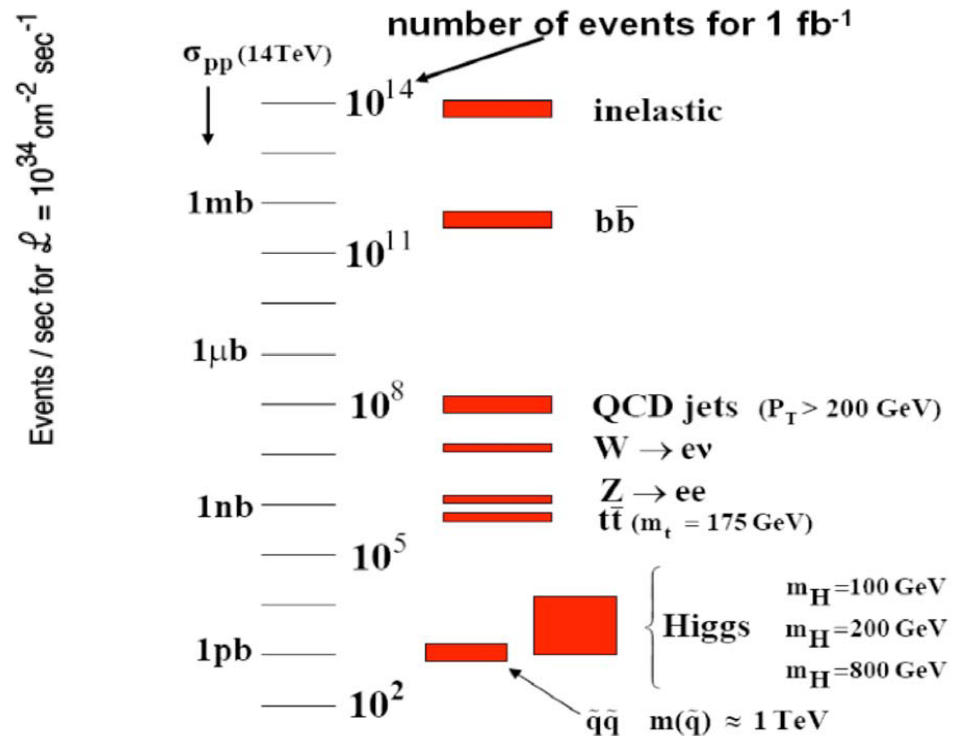
- Introduction
- How to install CMSSW
- Deployment of CMSSW on Grid sites
 - Performance
- Summary

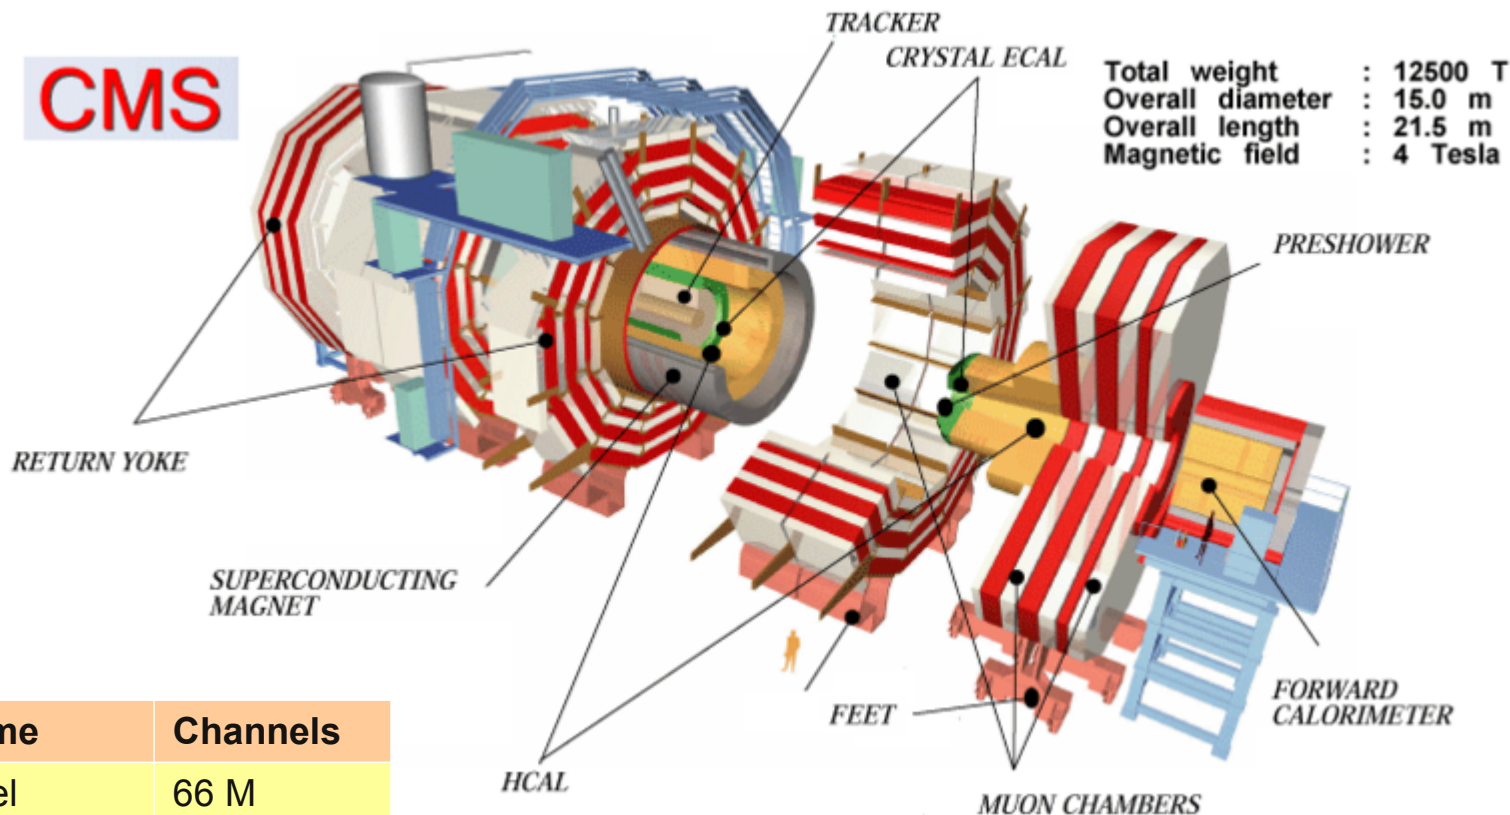


Motivation

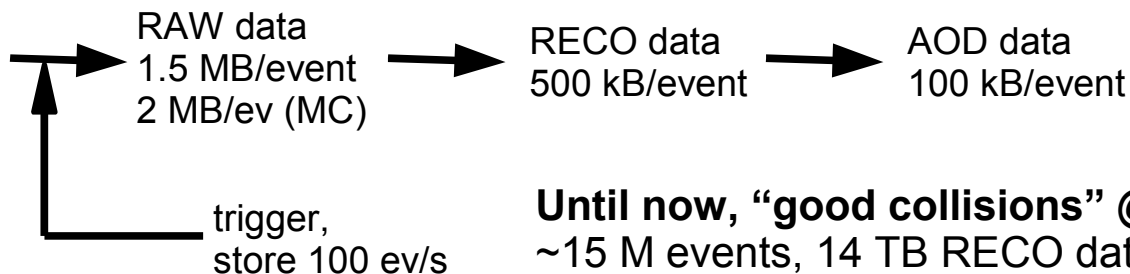


- Increase CM energy and luminosity
- Rediscover known particles, e.g. top
- Search for NP: Higgs, SUSY, ...





Name	Channels
Pixel	66 M
SiStrip	9.3 M
Muons	890 k
Preshower	140 k
ECAL	76 k
HCAL	9 k



**Until now, "good collisions" @7TeV:
~15 M events, 14 TB RECO data**


Data from: 2008 JINST 3 S08004, Table 9.1



CMS Centers around the world



More than 50 CMS centers, 22 countries



> Tier 0

- prompt reconstruction
- store RAW data and export to T1s

> Tier 1

- re-reconstruction
- long term storage of RAW data

> Tier 2

- MC production
- User analysis

> Tier 3

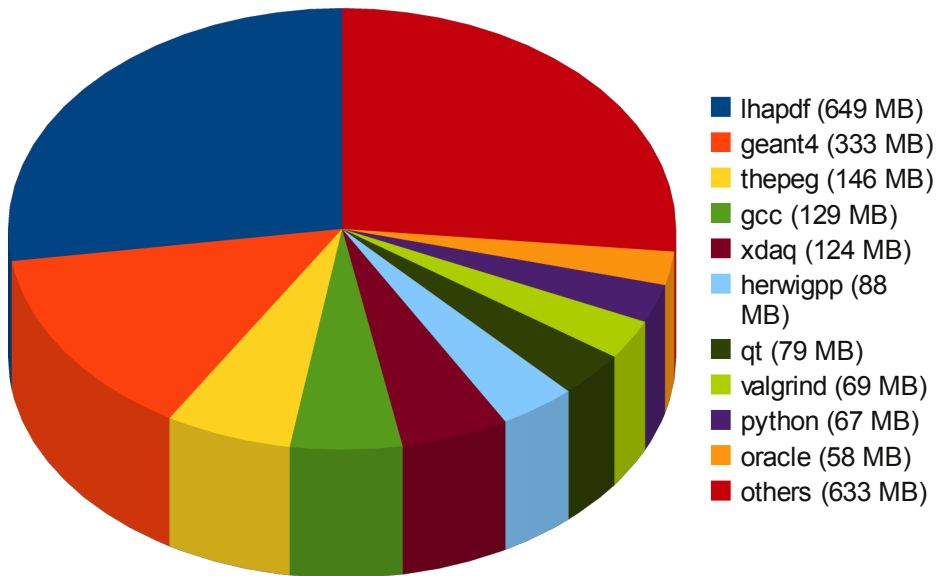
- User analysis

All centers need CMS software!





- Huge monolithic build
 - Contains:
MC, DQM, Reco, Analysis, Event Displays, ...
- Supported OS: Scientific Linux
- Lots of external dependencies:



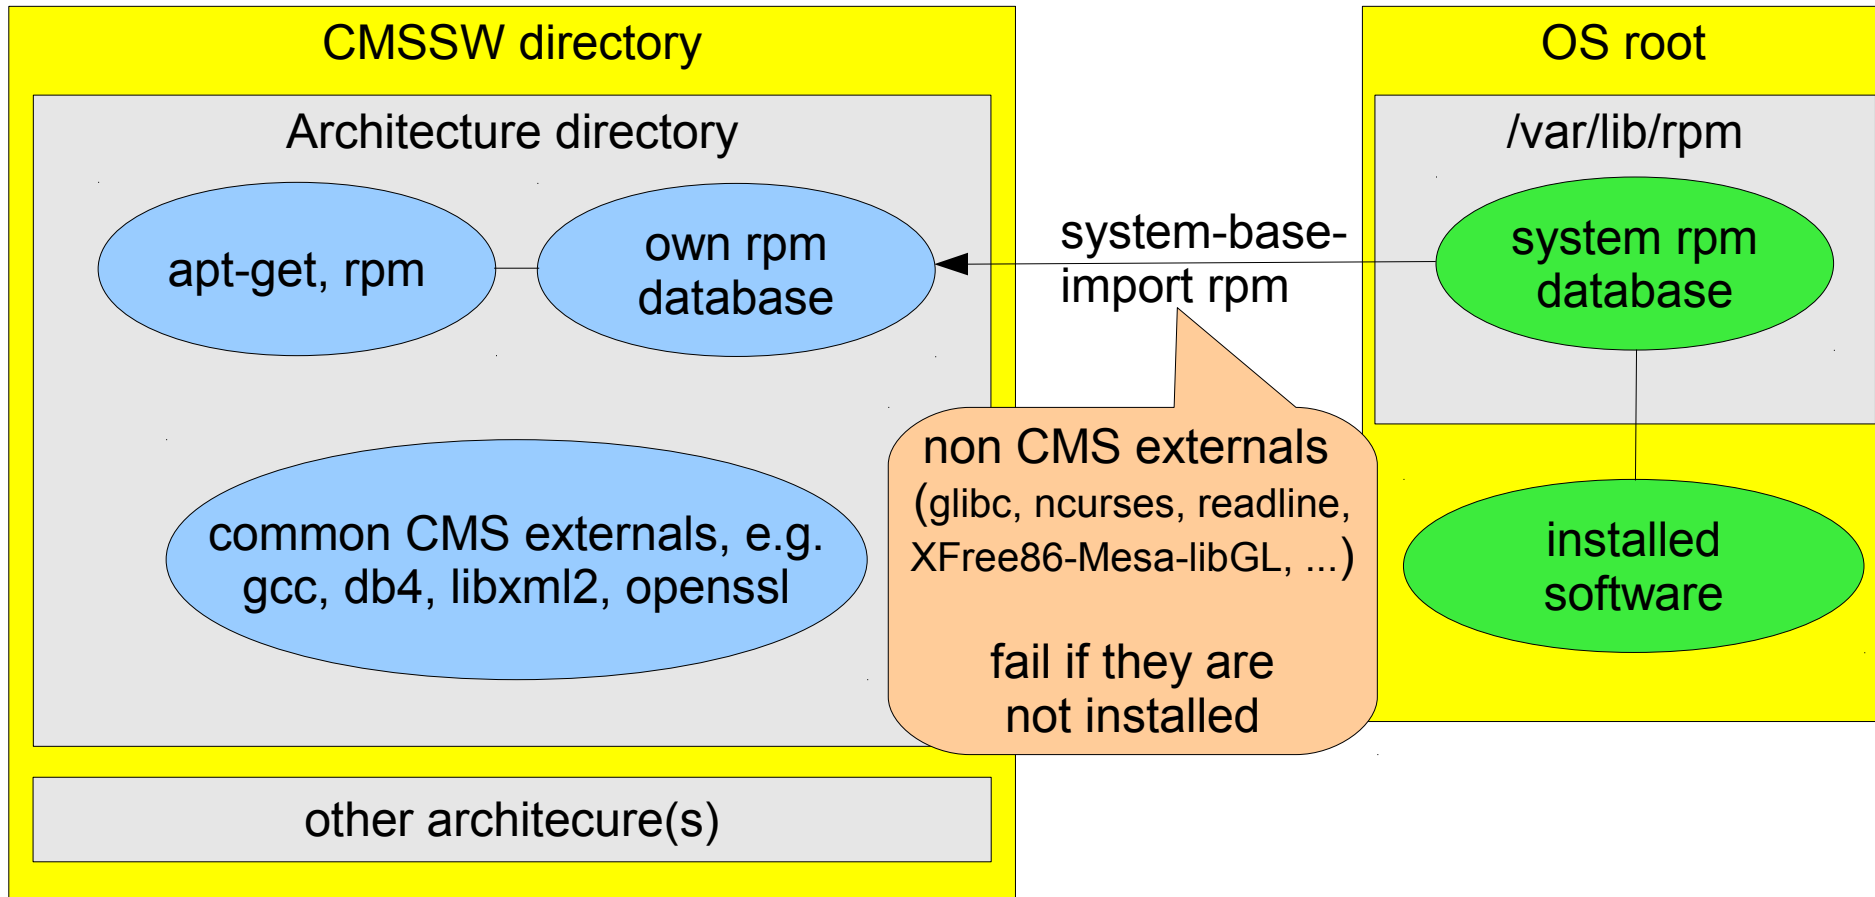
- The size of CMSSW (3_6_0)
 - 5.5 GB of files
 - 114,598 files and 4,197 symlinks in 11,115 directories
 - 4.8 million lines of code:
*.cc: 2.5m, *.h: 1m, *.py: 0.8m
 - 2.4 GB externals in 83 packages
- Different CMSSW versions can share packages
 - Disk size of a second release depends on how much is shared
- Too big for runtime install!



Installing CMSSW: step 1: Bootstrap



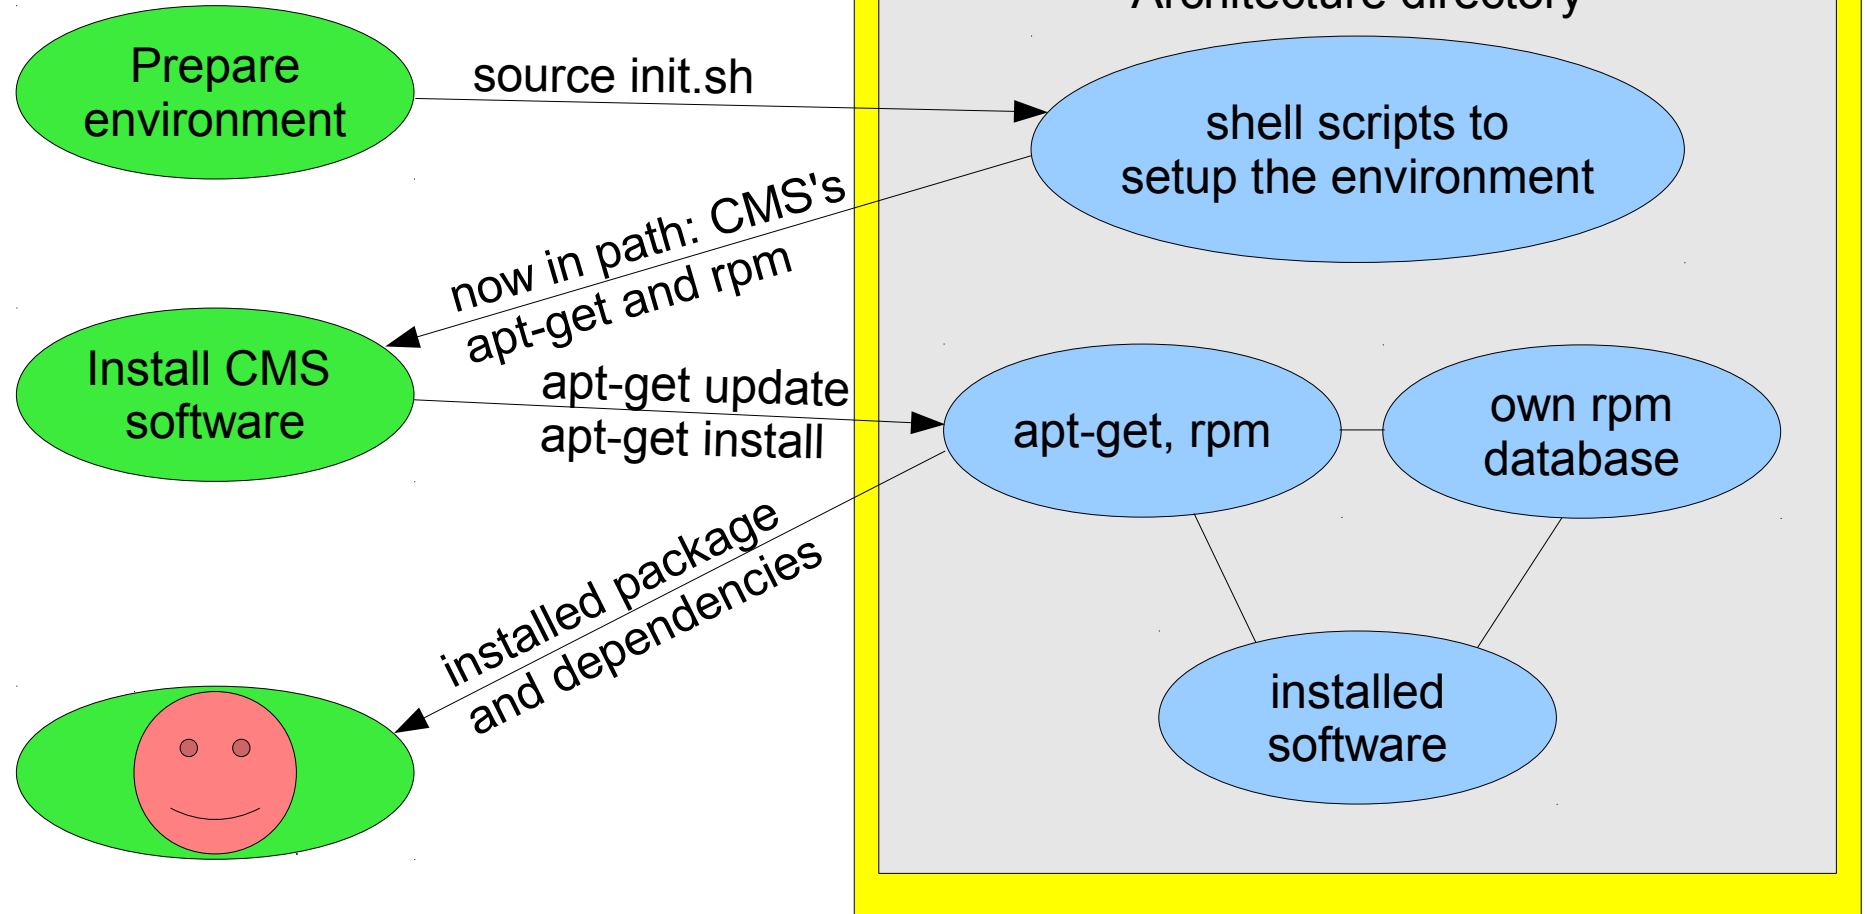
- Setup CMSSW in a user directory (no installation as root!)
- Needs to be done *once* per architecture (currently SL5, ia32, gcc434)



Installing CMSSW: step 2: Installing a Release



- Completely independent of system OS!



> CMSSW can be installed

- on your laptop / desktop



- on a file server in your local network



- on Grid sites



> CMSSW on the Grid

- central deployment
- deployment is not a site responsibility unless requested
- deployment team takes care of problems, contact site only if needed (e.g. permission errors)

> Two deployment teams (differences in Grid middleware)

- OSG – North and South America, direct installation on CE
- EGEE, ARC – Europe and Asia, installation via Grid job from a worker node

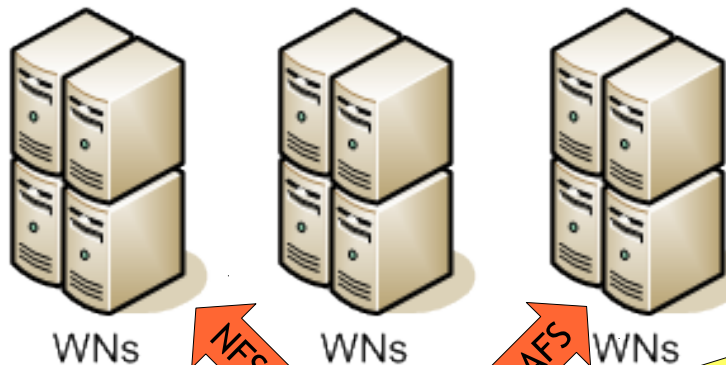

CMSSW Installation on a Grid Site (EGEE)



Usual configuration of one site with one CE

CMSSW mounted in `$VO_CMS_SW_DIR`
Can be read by any VO member
Write access requires VOMS Role=lcgadmin

CE info provider:
publishes installed
software (tags)



connected via
a network
file system



Fileserver
NFS, AFS, ...

CMSSW
installed on
file server

Gatekeeper:
• checks user certificate
• VOMS role dependent mapping

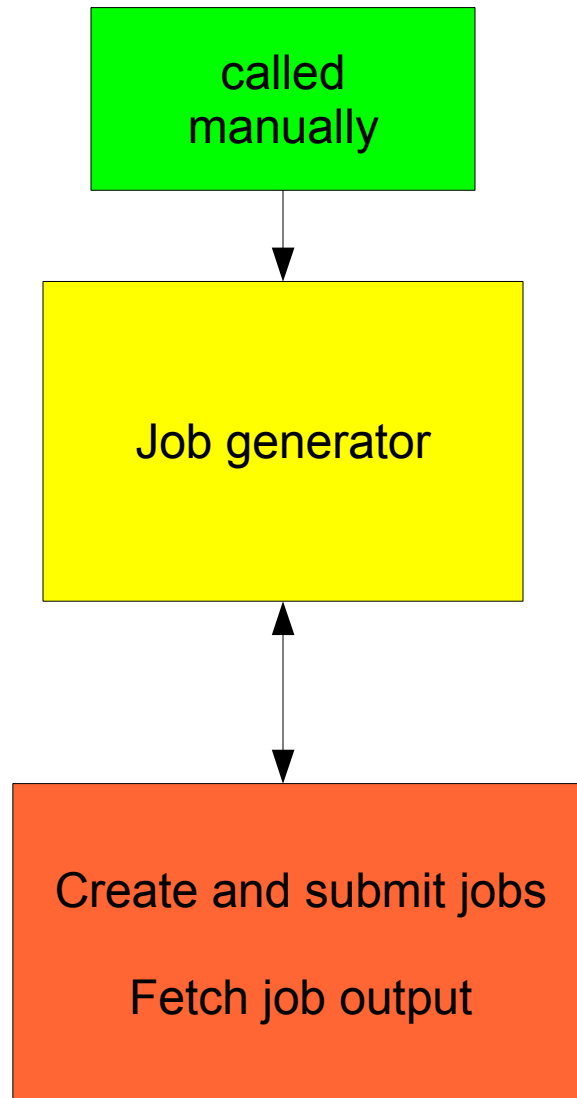




No interactive access to CE, installation of CMSSW via Grid job

- Grid job needs to be robust
 - Debugging very difficult
 - Log files and error messages must be as verbose as possible
- Deal with different site setups
 - Sites with multiple CEs
 - Shared file server?
 - Shared published software releases?
 - Operating System
- Work around most common problems
 - Different file systems, locking issues (move RPM database away)
 - Publishing tags fails (remove lock file)





Where to install?

List of CEs

called
manually

Job generator

Create and submit jobs
Fetch job output



➤ The CMS SiteDB

- CMS internal information
- site information, contact people
- available resources, list of CEs
- provides XML and JSON feeds
- *site admins can decide to install manually (no feed yet)*
- *release pinning: keeping certain deprecated releases (no feed yet)*

➤ List of all CEs also available via BDII

- might also include CEs which are not fully setup

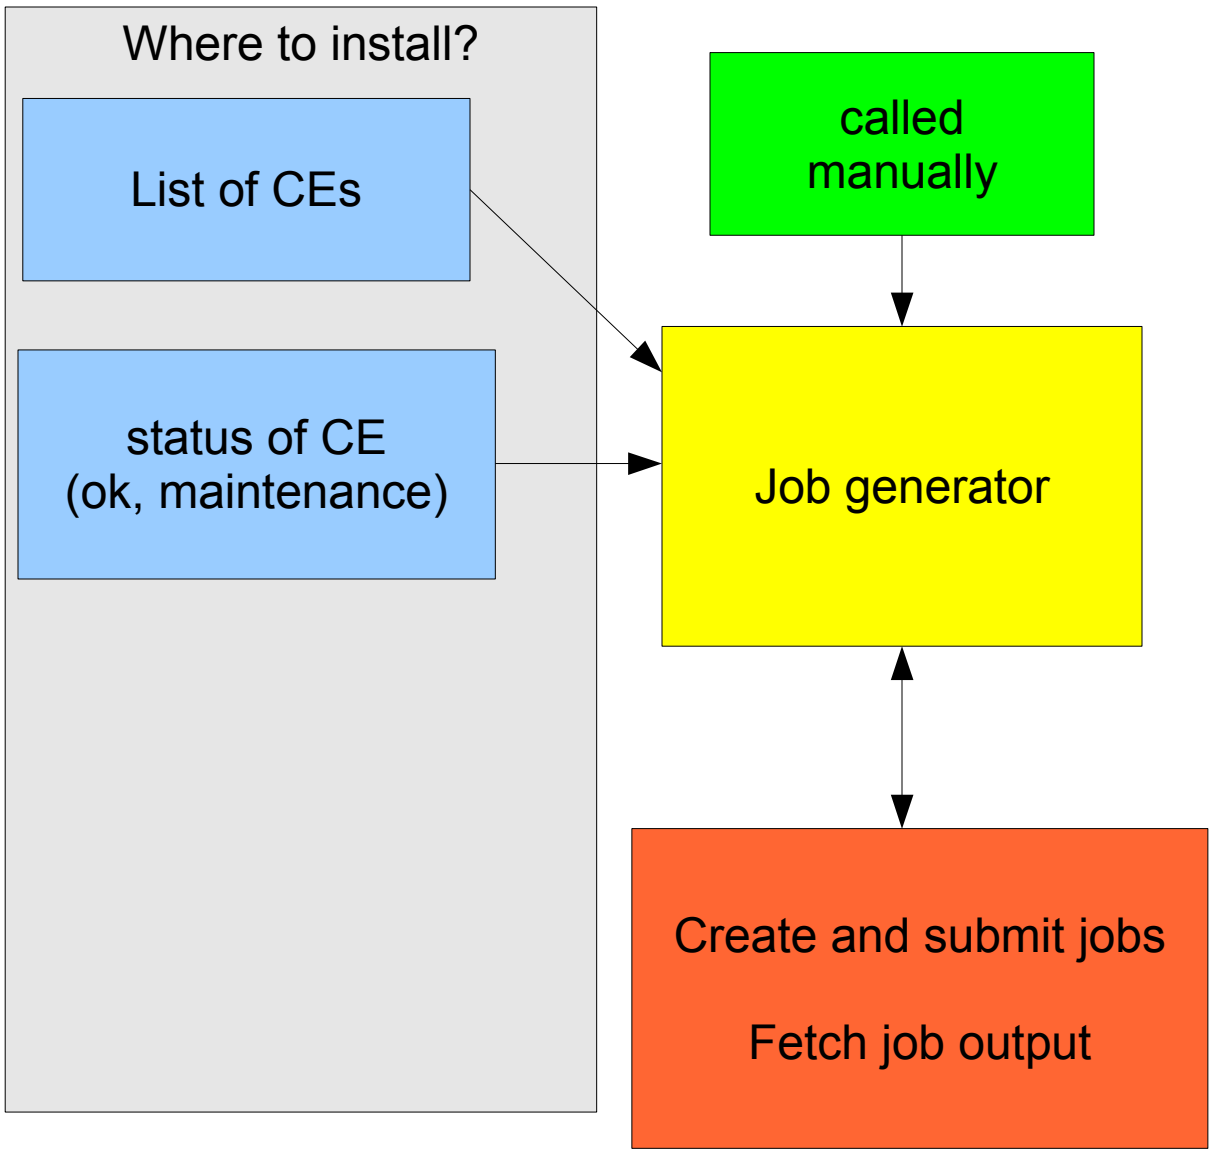
➤ Using a manual list

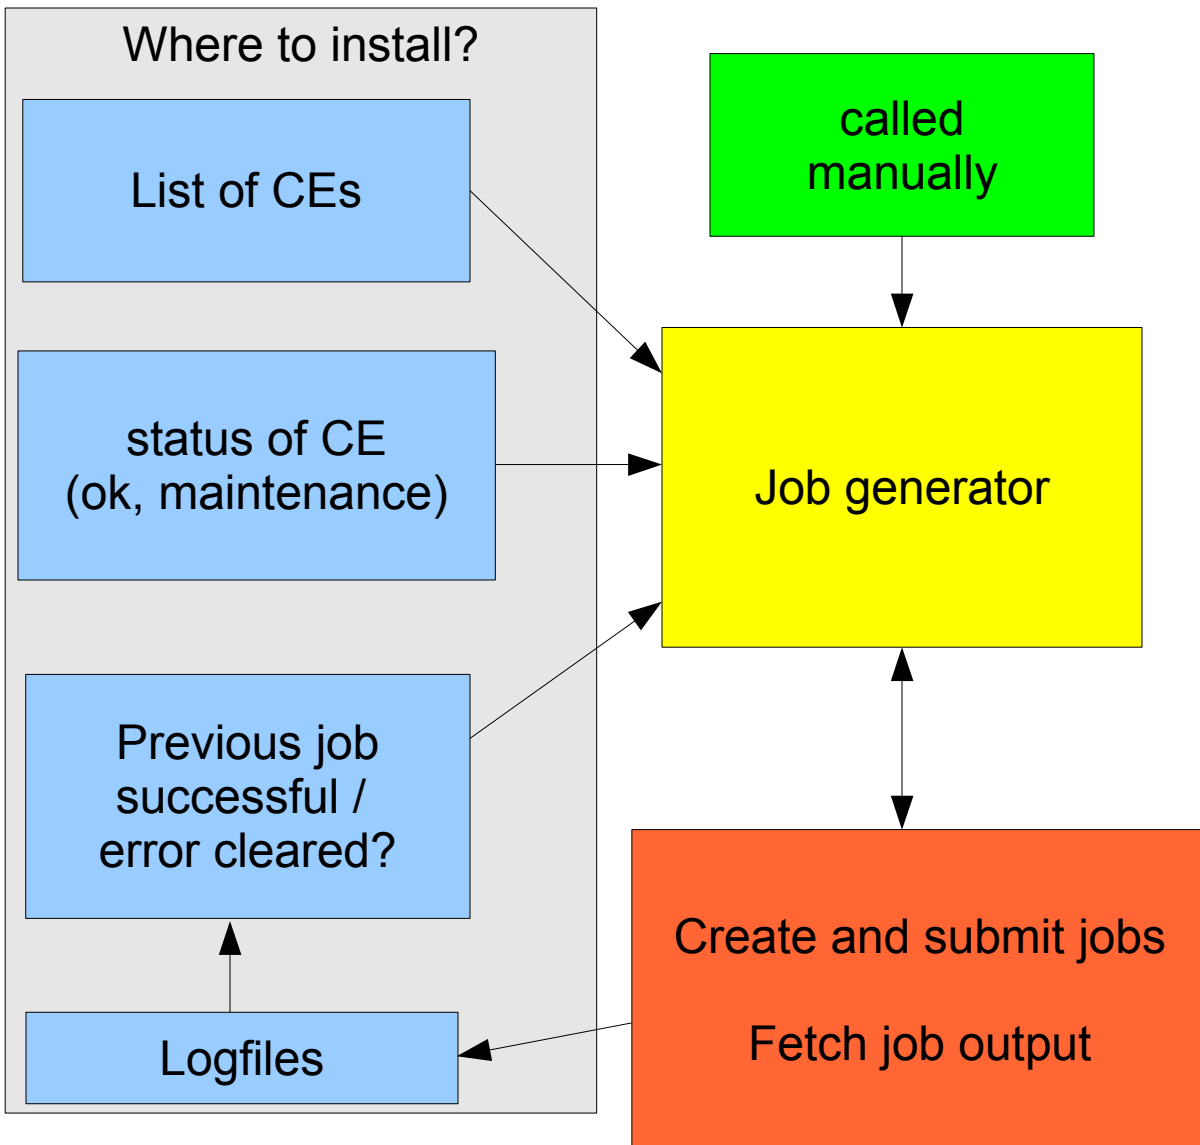
➤ Automatic detection of shared software areas

- sites with more than one CE might or might not share their software area
- ensures only one install job is sent per software area

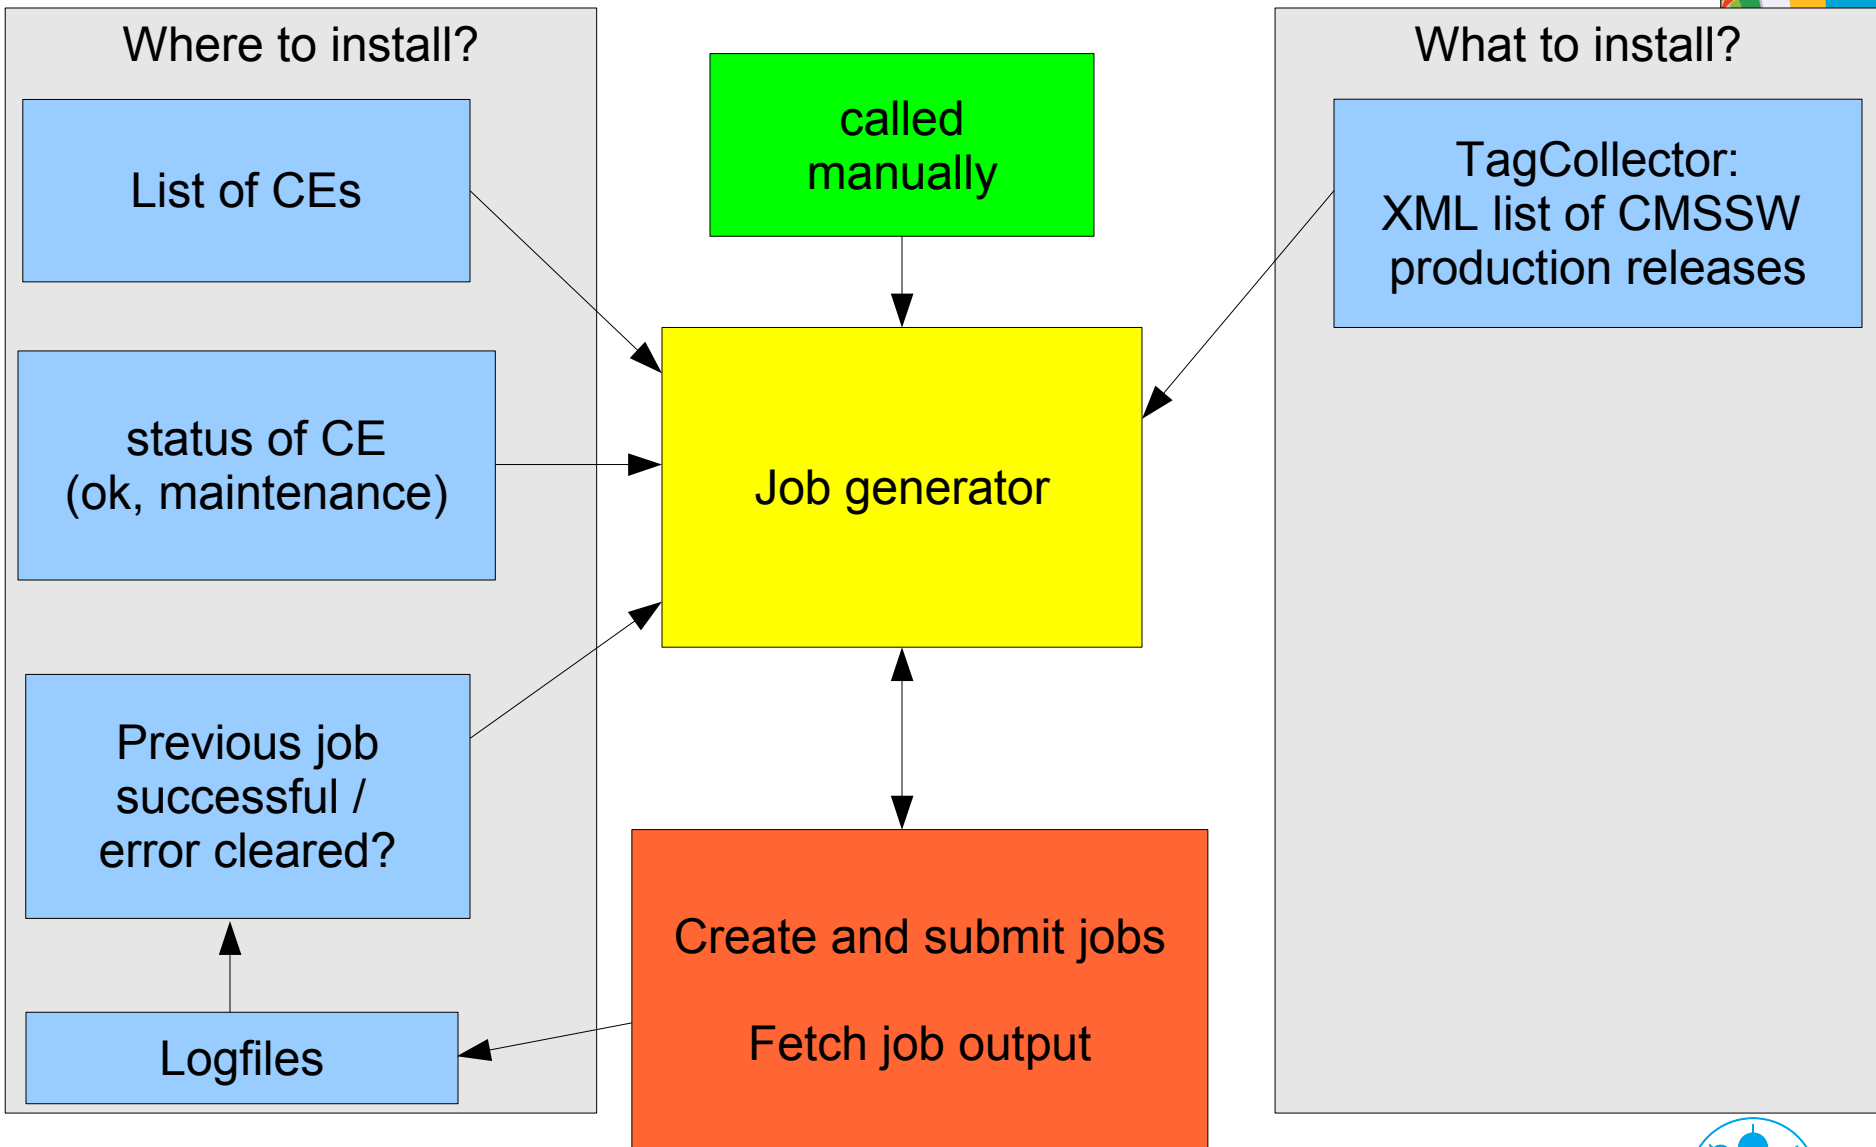
Edit Site Configuration			
CE(s) FQDN (s): comma separated	<input type="text" value="grid-ce5.desy.de, grid-ce4.desy.de"/>	Select Primary CE:	dropdown box
SE(s) FQDN (s): comma separated	<input type="text" value="dcache-se-cms.desy.de"/>	Select Primary SE:	dropdown box
GOCD ID:	<input type="text" value="35"/>	Install development CMSSW releases?: <input type="checkbox"/>	Site installs software manually?: <input type="checkbox"/>
		<input type="button" value="Submit"/> <input type="button" value="Cancel"/>	







Central Installation





Where to install?

List of CEs

called
manually

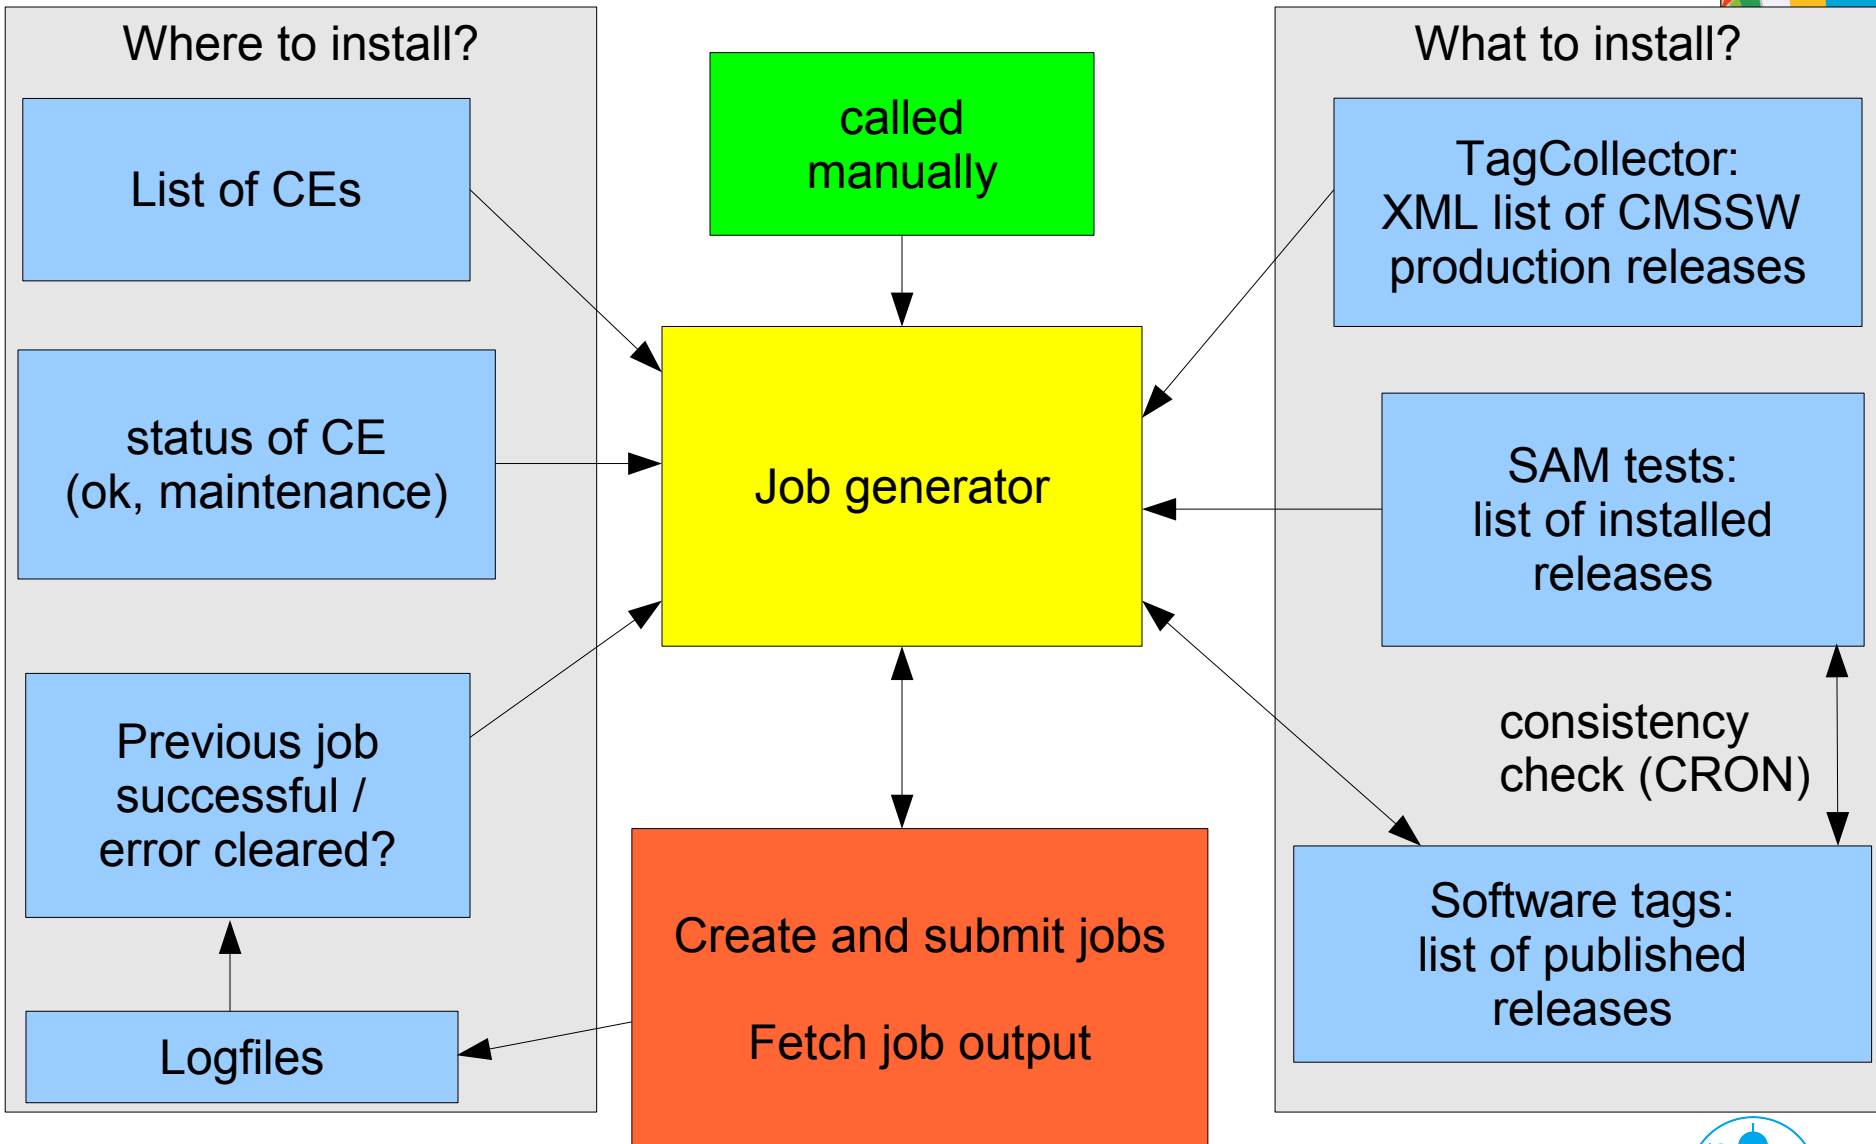
What to install?

TagCollector:
XML list of CMSSW
production releases

```
<projects>
<architecture name="slc5_ia32_gcc434">
  <project label="CMSSW_3_6_0" type="Production" state="Announced"/>
  <project label="CMSSW_3_6_0_patch2" type="Production" state="Announced"/>
  (...)
  <project label="CMSSW_3_4_0" type="Production" state="Announced"/>
</architecture>
<architecture name="slc4_ia32_gcc345">
  <project label="CMSSW_3_3_6" type="Production" state="Announced"/>
  <project label="CMSSW_3_3_6_patch6" type="Production" state="Announced"/>
  (...)
  <project label="CMSSW_2_2_6" type="Production" state="Announced"/>
</architecture>
</projects>
```



Central Installation



What to install?



➤ Service Availability Monitoring (SAM) tests are sent to all CEs every hour, results displayed on the CMS dashboard


➤ swinst test for deployment

- software installed correctly?
- releases missing?
- software tags correct?
- detailed information available, summary status is clickable
- sometimes a few hours old

➤ Software tags

- tags set directly after installation, no delay
- Nordugrid (ARC middleware): deployment jobs cannot set tags

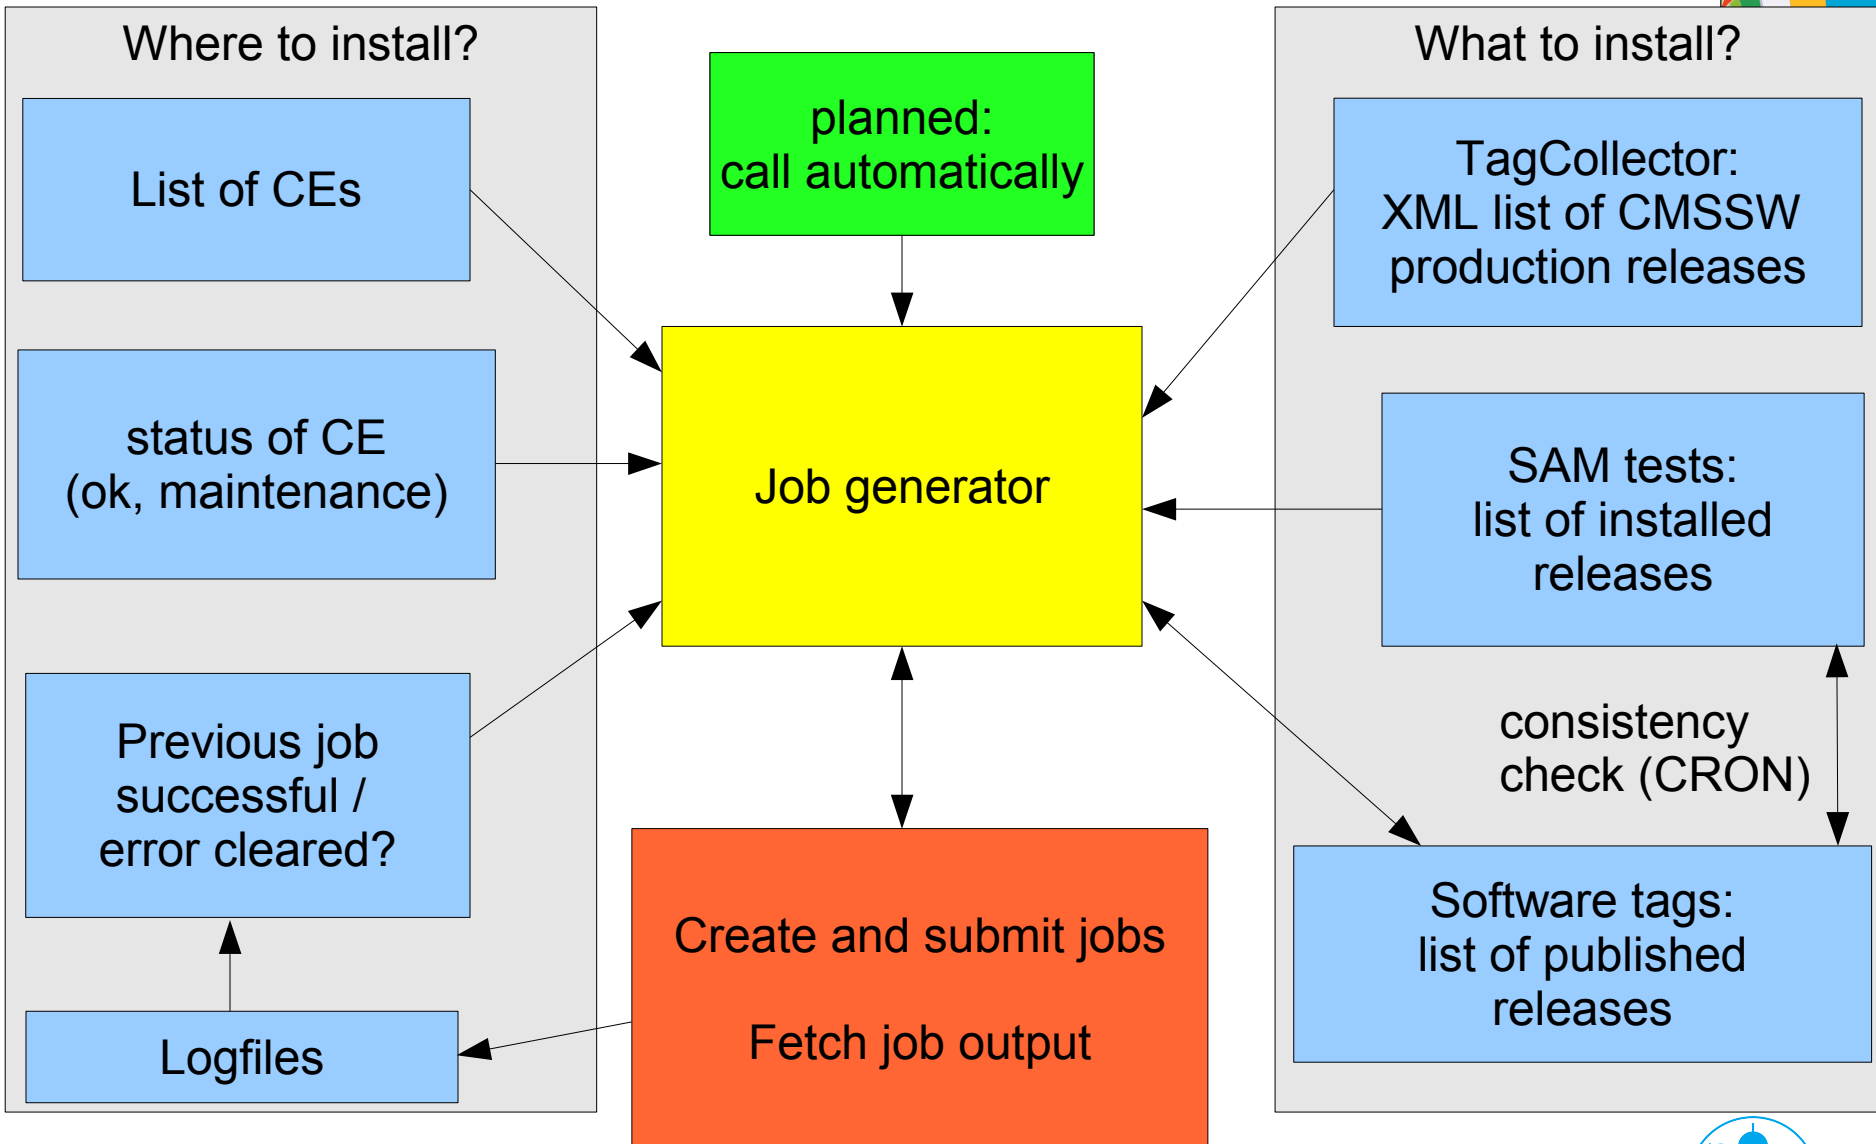
➤ Using a combination of SAM tests and software tags to determine installed releases



Sitename	Service Type	Service Name	mc	sft-job	swinst	squid	frontier	basic	prod	analysis
T1_DE_KIT	CE	ce-2-fzk.gridka.de	ok	ok	ok	ok	ok	ok	ok	ok
		ce-3-fzk.gridka.de	ok	error	ok	ok	ok	ok	ok	ok
		ce-4-fzk.gridka.de	ok	ok	ok	ok	ok	ok	ok	ok
T1_ES_PIC	CE	ce05.pic.es	ok	ok	ok	ok	ok	ok	ok	ok
		ce06.pic.es	ok	ok	ok	ok	ok	ok	ok	ok
		ce07.pic.es	ok	ok	ok	ok	ok	ok	ok	ok
T1_FR_CCIN2P3	CE	ccicgcell04.in2p3.fr	ok	ok	ok	ok	ok	ok	ok	ok
		ccicgcell08.in2p3.fr	ok	ok	ok	ok	ok	ok	ok	ok
T1_IT_CNAF	CE	ce04-lcg.cr.cnaf.infn.it	ok	ok	ok	ok	ok	ok	ok	ok
		ce05-lcg.cr.cnaf.infn.it	ok	ok	ok	ok	ok	ok	ok	ok
		ce06-lcg.cr.cnaf.infn.it	ok	ok	ok	ok	ok	ok	ok	ok
T1_TW_ASGC	CE	lcg00125.grid.sinica.edu.tw	ok	ok	ok	ok	ok	ok	ok	ok
		w-ce01.grid.sinica.edu.tw	ok	ok	ok	ok	ok	ok	ok	ok
		w-ce02.grid.sinica.edu.tw	ok	ok	ok	ok	ok	ok	ok	ok
T1_UK_RAL	CE	w-ce03.grid.sinica.edu.tw	ok	ok	ok	ok	ok	ok	ok	ok
		lcgce06.gridpp.rl.ac.uk	ok	ok	ok	ok	ok	ok	ok	ok
T1_US_FNAL	CE	lcgce07.gridpp.rl.ac.uk	ok	ok	ok	ok	ok	ok	ok	ok
		cmsosgce.fnal.gov	ok	ok	ok	ok	ok	ok	ok	ok
		cmsosgce2.fnal.gov	ok	ok	ok	ok	ok	ok	ok	ok
T2_AT_Vienna	CE	cmsosgce4.fnal.gov	ok	ok	ok	ok	ok	ok	ok	ok
		hephygr.oeaw.ac.at	ok	ok	ok	ok	ok	ok	ok	note
T2_BE_IHHE	CE	lcgce.oeaw.ac.at	ok	ok	ok	ok	ok	ok	ok	ok
		cream01.ilhe.ac.be	ok	ok	ok	ok	ok	ok	ok	ok
T2_BE_UCL	CE	gridce.ilhe.ac.be	ok	ok	warn	ok	ok	ok	ok	note
		lncgrid.cism.ucl.ac.be	ok	ok	ok	ok	ok	ok	ok	note

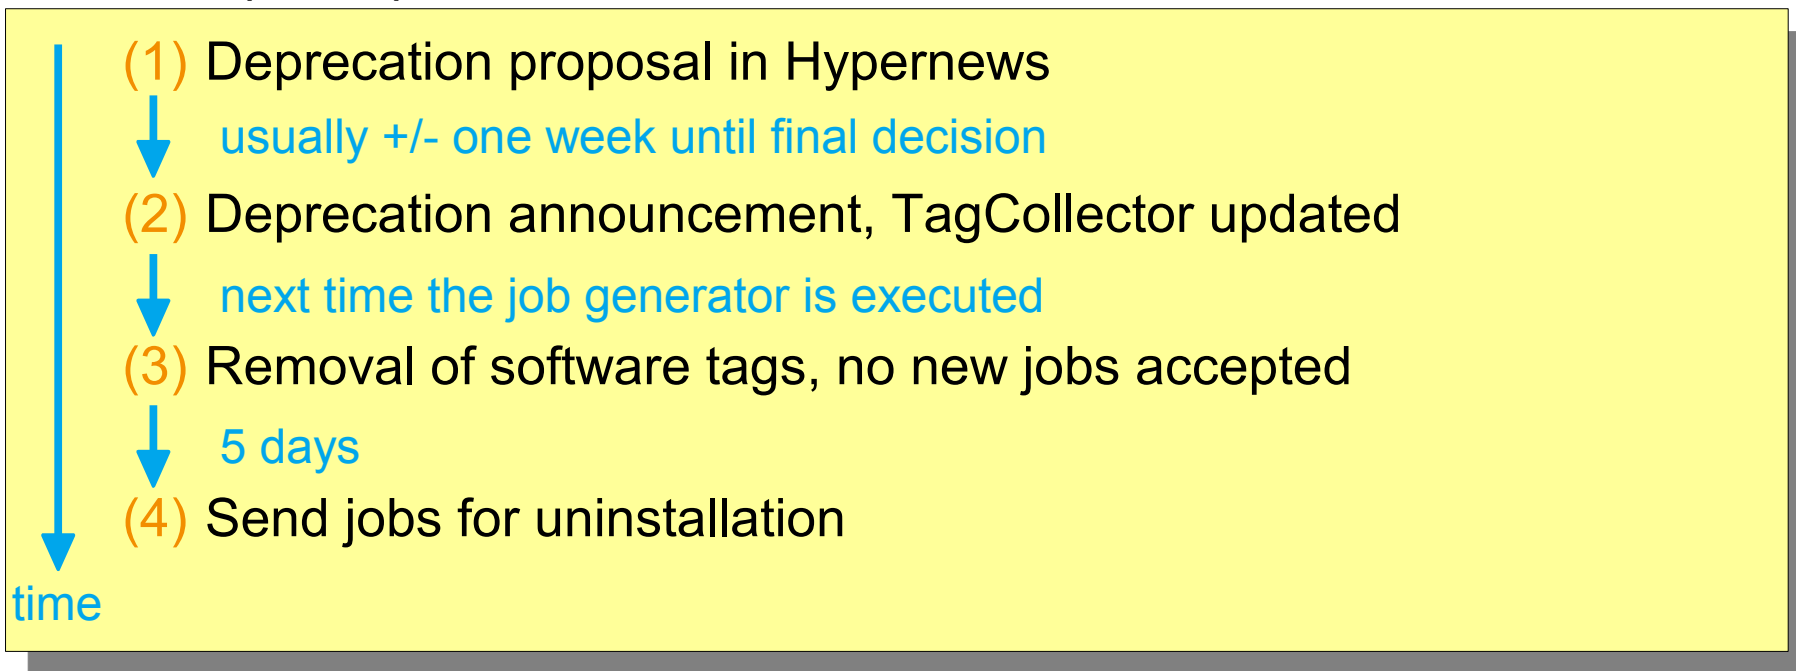


Central Installation





- Removal of old releases required
 - New releases contain better algorithms and bug fixes
 - Disk space, rpm limitations



- Well-defined way to remove old releases
 - Running and/or queued jobs will not fail due to removal of software





- Well documented Perl code, object oriented
 - previous deployment software also Perl, personal preference
- 9 classes, 5700 lines of code

LcgCE
Compute Element specific functions

JDL
Grid job submission

CMSSWInstall
Manage CMSSW installations

- Installation can be tested locally:
 - create `CMSSWInstall` object, call `$obj->install("CMS_version")`
- Code reuse
 - Dashboard/SiteDB comparison, user script submission to all CEs, ...
 - Monitoring web site





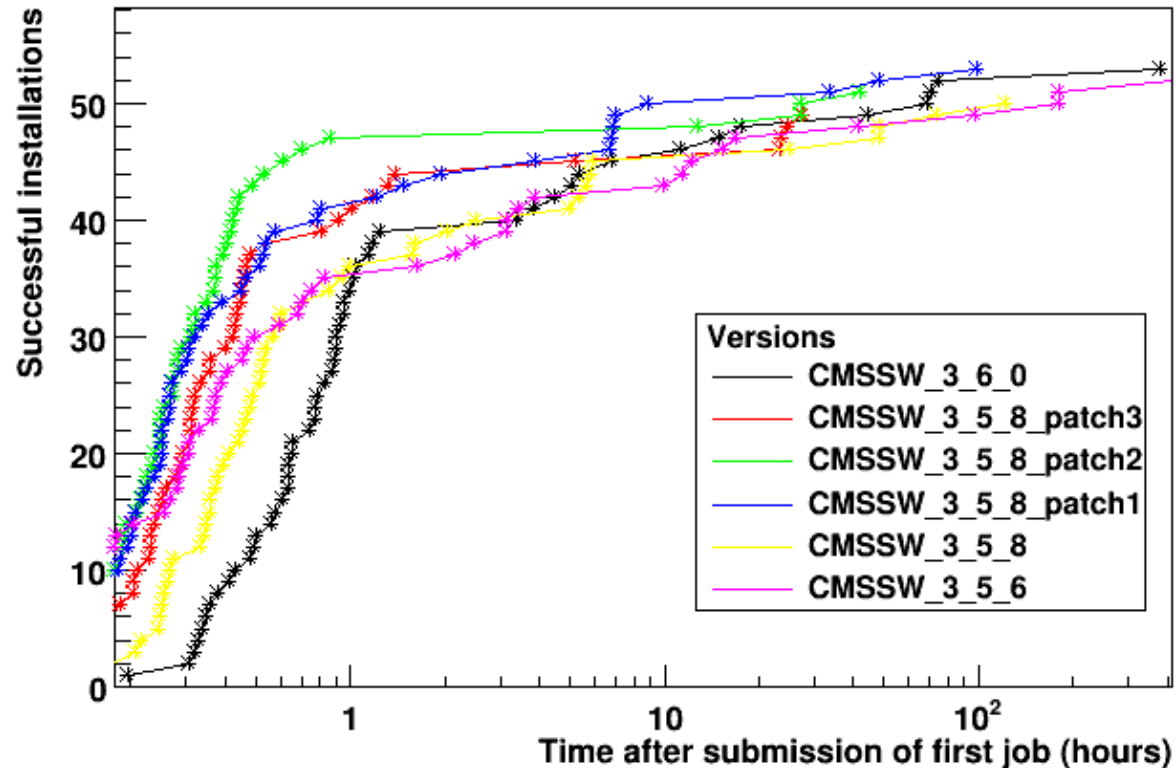
➤ Very good performance

- ~40 installs within 1 hour
- clearing the tail causes a lot of work

➤ Reasons for delayed deployment

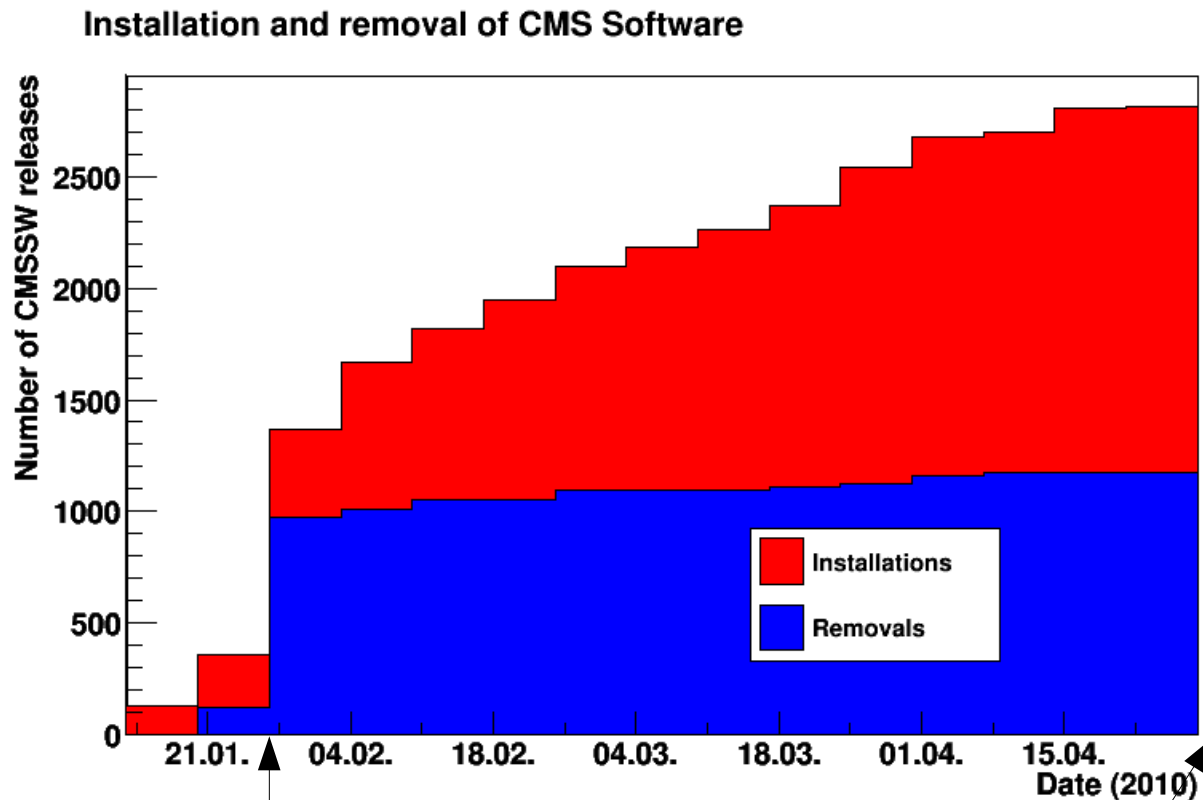
- CE is down/maintenance
- jobs need to wait in the queue
- slow network connection
- slow file system - sometimes causes more than a factor of 100 in speed!
- other technical problems, savannah tickets

Performance of the CMSSW installation





- Total numbers since January
 - > 100 different releases
 - > 80 CEs
 - > 1600 successful jobs
 - > 3100 releases installed
- Average job:
 - queues for 20 mins
 - runs 80 mins
 - installs 1.9 releases



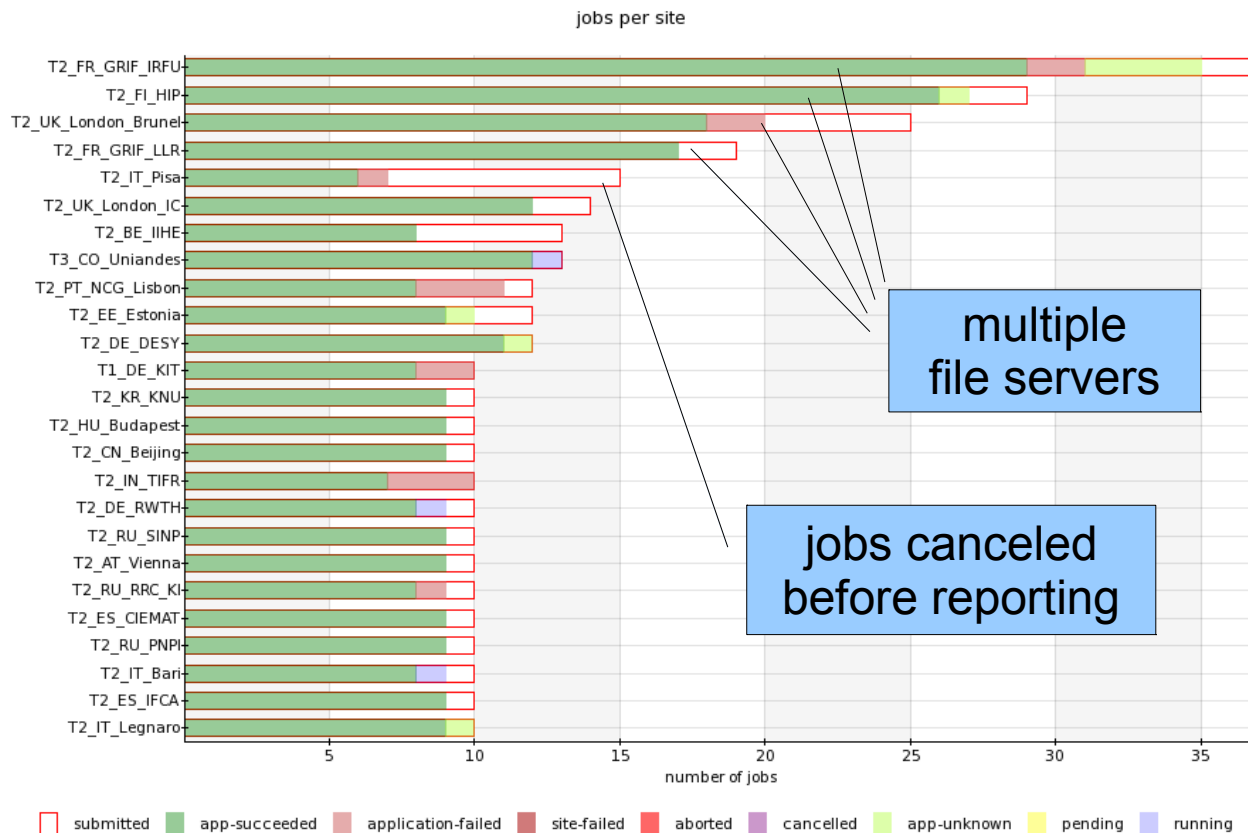
lots of old releases deprecated at the same time

43 production releases installed (May 7)





- CMS standard
- Base/generic job information (submitted, running, success / error)



jobs since April 2010





➤ Static html page

- based on SAM tests and log files
- similar to SAM test overview
- missing releases wrt. OS version
- shows times and dates of installation
- needs to be created manually and copied to a web server

Tier	Site	CE	OS	Time	3	3	3	3	3	3	3	3
					6	6	6	6	6	6	6	
					0	0	0	0	0	0	0	0
					patch2	patch1	pre5io	pre5	pre3	patch3	patch2	
2	Brunel	dgc-grid-44.brunel.ac.uk	SL 4.7	0.4 h								
2	CIEMAT	lcg02.ciemat.es	CERN SLC 5.4	0.8 h	X 04.05. 20:45 -8/16	X 04.05. 14:12 -2/3	X 16.04. 18:18 16/8	X 12.04. 12:50 5/21	X 12.04. 12:50 5/21	X 24.03. 21:12 5/19	X 03.05. 17:38 5/3	X 30.04. 16:49 0/2
2	CSCS	ce01.lcg.cscs.ch CMSSW_3_5_8_patch3 CMSSW_3_6_0_patch2	SL 5.3	0.8 h			X 16.04. 23:03 4/313	X 23.04. 21:19 4/299	X 23.04. 21:19 4/299			X 30.04. 16:56 5/17
2	DESY	grid-ce4.desy.de	SL 5.3	0.4 h	X	X	X			X 24.03. 18:29 4/35	X	X
2	DESY	grid-ce5.desy.de	SL 5.3	0.8 h	X 04.05. 20:08 -4/2	X 04.05. 14:03 1/2	X 16.04. 18:03 3/15			X	X 03.05. 17:26 4/2	X 30.04. 16:41 1/2
2	Estonia	oberon.hep.kbfi.ee	SL 5.3	0.5 h	X 04.05. 20:37 -7/3	X 04.05. 14:13 -5/4	X 19.04. 20:26 1/35				X 03.05. 17:36 1/3	X 30.04. 16:47 -6/3
2	IPHC	sbgee1.in2p3.fr	SL 5.4	0 h	X 04.05. 20:35 -10/3	X 04.05. 14:12 -6/3	X 16.04. 18:19 0/24				X 03.05. 17:34 0/3	X 30.04. 16:46 -8/3
2	ITEP	ceglite.itep.ru	CERN SLC 5.5	0.5 h	X	X	X				X	X
tier unknown	site unknown	ce3.itep.ru	CERN SLC 5.5	0.8 h	X 04.05. 20:18 4/4	X 04.05. 14:06 4/4	X 16.04. 18:05 4/15				X 03.05. 17:43 20/3	X 30.04. 16:42 4/3
2	ITEP	ceitep.itep.ru	CERN SLC 4.8	0.8 h	X	X	X				X	X
2	JINR	lcgce01.jinr.ru CMSSW_3_5_8_patch3 CMSSW_3_6_0_patch2	SL 5.4	0.8 h			X 17.04. 08:41 2/22					X 30.04. 16:49 -2/3
2	JINR	lcgce02.jinr.ru CMSSW_3_5_8_patch3 CMSSW_3_6_0_patch2	SL 5.4	1.4 h			X					X





➤ Submission monitoring

- dynamic pages
- up-to-date deployment performance plots
- not a “duplicated dashboard”
- Easy browsing through error log files

A screenshot of a web browser window titled 'CMSSW Deployment Monitoring'. The address bar shows 'http://127.0.0.1/cgi-bin'. The page content includes a navigation bar with four buttons: 'Overview' (blue), 'CEs in Error' (blue), 'CEs with missing Releases' (blue), and 'Running Jobs' (red). Below the navigation bar is a table with four columns: 'CE', 'Submitted', 'Task', and 'Job ID'.

CE	Submitted	Task	Job ID
dc2-grid-65.brunel.ac.uk	2010-05-05 09:47:44	Install: CMSSW_3_6_0_patch2	https://...
grid-ce-01.ba.infn.it	2010-05-04 20:23:16	Install: CMSSW_3_6_0_patch2	https://...

➤ Currently in very early stage

- running on local webserver, using a local copy of the log files

➤ Ideas for a faster deployment process

- connect with site downtime calendar
- display link to savannah page in error log pages
- clear error states and resubmit jobs
- **security becomes an important issue once doing more than monitoring!**





S

CMSSW Deployment Monitoring

http://127.0.0.1/cgi-bin/test.pl?stoppedces=1

CMSSW Deployment Status

Overview | **CEs in Error** | **CEs with missing Releases** | **Running Jobs**

dc2-grid-65.brunel.ac.uk - grid executable didn't return success! - [Open Savannah ticket](#)

Systeminfo:

[\[Start timestamp\]](#) [\[User information\]](#) [\[OS information\]](#) [\[umask\]](#) [\[Grid Proxy Information\]](#)
[\[Environment variables\]](#) [\[Mounted filesystems\]](#) [\[Free space on filesystems\]](#) [\[Directory listings\]](#) [\[end sysinfo at:\]](#)

Errorlog:

```
Executing external command failed!  
Command: lcg-tags --ce dc2-grid-65.brunel.ac.uk --vo cms --add --tags VO-cms-CMSSW_3_6_0_patch2-  
processing-install-2010/05/10_12:32:31 > /scratch/174400.dc2-grid-65.brunel.ac.uk/MiscToolsRun-out-  
FrNNxHhw 2> /scratch/174400.dc2-grid-65.brunel.ac.uk/MiscToolsRun-err-nWUCp5Ac  
Output (stdout) so far:  
  
Output (err) so far:  
lcg-tags: error: lock file already present at gsiftp://dc2-grid-65.brunel.ac.uk/opt/edg/var/info/cms/  
lock.  
lcg-tags: error: cannot create remote lockfile.  
Sorry, I can't continue because of the error at MiscTools.pm line 139, <$fh> chunk 1.  
  
Stdout:  
Creating new CMSSWInstall object  
Clearing RPM database cache...  
Cleaning repository...  
Executing apt-get --assume-yes clean all...  
Installing CMSSW version CMSSW_3_6_0_patch2...
```



C



lc

Jobs

b ID

s://...

s://...





➤ Submission monitoring

- dynamic pages
- up-to-date deployment performance plots
- not a “duplicated dashboard”
- Easy browsing through error log files

The screenshot shows a web browser window titled 'CMSSW Deployment Monitoring'. The address bar shows 'http://127.0.0.1/cgi-bin'. The page content includes a navigation bar with four tabs: 'Overview' (selected), 'CEs in Error', 'CEs with missing Releases', and 'Running Jobs'. Below the navigation bar is a table with the following data:

CE	Submitted	Task	Job ID
dc2-grid-65.brunel.ac.uk	2010-05-05 09:47:44	Install: CMSSW_3_6_0_patch2	https://...
grid-ce-01.ba.infn.it	2010-05-04 20:23:16	Install: CMSSW_3_6_0_patch2	https://...

➤ Currently in very early stage

- running on local webserver, using a local copy of the log files

➤ Ideas for a faster deployment process

- connect with site downtime calendar
- display link to savannah page in error log pages
- clear error states and resubmit jobs

▪ **security becomes an important issue once doing more than monitoring!**





- Deployment system works very well
 - Works around most common problems, thus high success rate
- Effort has been put mainly into automated job creation
 - just one command needed to install/remove releases on all CEs
- Following up problems can be time consuming

- Easier error handling in preparation
 - Use browser instead of shell
- More automated submission
- Better integration with SiteDB
 - No more manual list of CEs
 - Pinning of deprecated releases

