GRIF : The Challenge of a Distributed Site

Michel Jouvin
LAL, Orsay
jouvin@lal.in2p3.fr
http://grif.fr

October 8, 2008
GDB, CERN

Agenda

- GRIF project and resources
- Human challenges
- Importance of tools
  - Collaborative tools
  - Management tool (Quattor)
- Relation with experiments
- Conclusions

Why GRIF ?

- Goal : build a large EGEE/LCG node to support both central and end-user analysis for all LHC experiments
  - 80% of French LHC scientists in Paris region
  - 5 HEP labs in Paris region (IN2P3 + CEA/IRFU)
  - None of the labs large enough to do something alone
  - Avoid competition for funding and manpower in the same region
- Leverage on LAL and IRFU expertise since 2001
  - Other labs have no grid experience
  - LAL has experience with Quattor, a potentially suitable management tool
- Remain a resource opened to other communities
  - LAL and IRFU (DAPNIA) are significant contributors to biomed and ESR
  - LAL involved in NA4 (Cal) effort to help new users
  - Mandatory in French context to obtain regional funding

GRIF Project

- Grille au service de la Recherche en Ile de France
  - Started beginning 2005 as a joint effort by 5 labs
  - Early decision to create one site to provide consolidated resources and lower management load
  - Hardware distributed on each participating site
    - Avoid infrastructure problems, required to get local funding from each University or CEA.
  - Build a unique technical team made of volunteers from each lab
    - No possibility to hire anybody new
  - No formal structure created : too complex to start, rely on existing labs and their usual relationships
    - All procurements made independently by each lab, even though strongly coordinated (in particular contact with vendors)
- Scientific chairman : J.P. Meyer (CEA/IRFU)
- Technical coordinator : M. Jouvin (IN2P3/LAL)

GRIF Resources

- Initially 80% LHC, 20% others…
- Growth quicker than expected
  - Having a common project helped each partner to mobilize and consolidate local resources
    - One more lab (APC, astro-particles) joined 1 year ago
  - GRIF became attractive for a newcomer : Institut des Systèmes Complexes IdF
    - Important needs, no manpower but funding, decided to buy resources and put them in GRIF (with an equivalent guaranteed share)
- GRIF is probably one of the larger T2 site
  - CPU : 5 MSI2K (Clovertown 2.33 Ghz ~ 2 kSI2K/coeur)
    - Initial prevision for 2008 : 3 MSI2K
    - 3 MSI2K “pledged” (but not used) for ISC and benefiting to LCG
  - Disk : 500 TB
  - Network : 10 Gb/s inter-site, 10 Gb/s to CCIN2P3
    - GRIFOPN based on black fibres given by RENATER to support LCG
    - Demonstrated 250 MB/s sustained between CCIN2P3 et GRIF/LAL (1 SE)

Grid Core Services

- 1 top BDII (topbdii.grif.fr) used by many French sites
- VOMS
  - Used by 15-20 VOs, mostly small ones LFC Server
- LFC
  - Both a local server for some VOs like Atlas and the central server for VOs « supported » by GRIF
- MyProxy
- WMS
  - Currently 1 WMS + 1 LB + 1WMS/LB (gLite 3.0)
  - Heaviliy used by local users and some non-HEP VOs doing large productions
    - 1 person from VO astro submitting large number of DAG jobs
  - Moving to 1 LB shared by 2 or more WMS in different subsites (gLite 3.1)
    - Interested by recipes or future plans for LB replication
- GRIF chosen as 1 of the 4 EGEE « seed resource » providers

Human Challenge

- Build a unique technical team with people not at the same location and with other activities
  - Was considered since the beginning as critical for success
  - Could not expect any new manpower centralized on 1 location
  - Wanted to increase global grid expertise and avoid a split between grid and non grid people
- Technical committee opened to everybody motivated
  - 20 people representing 10 FTE
  - A core group of (4-5 people) grid experts
  - No formal or hierarchical structure
  - 1 F2F meeting/month + every day communication by email
    - Looking at setting up an IRC channel
- Take the time necessary to build consensus
  - Increase the chance of long-term choices
  - Increase motivation and implication of every body

Importance of Tools

- Tools helping a successful human organization but not a work-around to a malfunctioning one
  - In GRIF, importance reinforced by geographical distribution
- Collaborative tools for communication and traceability
  - Allow every day discussion with people availability changing every day
  - Involve every body in producing/maintaining documentation
  - Traceability of actions
  - Follow-up and planning of actions
  - Information exchange : meeting minutes considered important
- Management tools allowing distributed but consistent administration of the whole site
  - Avoid as much as possible duplication of effort
  - Allow non-experts to participate and increase their expertise
  - Consistency check of any change before deployment and ability of rolling back any change easily

Trac

- Combine a wiki, an issue tracker and a SVN client
  - Much more through additionnal plug-ins
- Corner stone of collaborative tools at GRIF
  - Wiki : easy contribution by many people to the internal documentation, easy to use
    - Simple syntax, no specific client required
    - Documentation is used and maintained at the same time
  - Cross-reference between any kind of information
    - E.g. : 1 ticket may reference documentation or a configuration change, a wiki page may reference a list of actions (tickets)…
  - SVN client gives access to configuration changes and history
- Access restricted to GRIF members
  - Allows to put any kind of information, even sensible one

Trac : Example ...

Doc/FAQ/Restart - Exploitation Informatique - Trac - Opera

Fichier Editer Afficher Signets Widgets Flux RSS Discussion Outils Aide

Nouvel onglet

https://trac.lal.in2p3.fr/Exploitation/wiki/Doc/FAQ/Restart

CNRS (FR) Google

Catégorie	Problème	Analyse	Action	Ticket/ Resp.
Problèmes électriques				
	HP1 disjoncte	Trop grosse charge éléctrique	Redémarrer les satellites du PDU 1 par 1	GP
Réseau				
	DNS lal.in2p3.fr + grif.fr inaccessible	CC IN2P3 n'était plus secondaire pour ces zones	Redémarrage de quattorsrv	~~#38~~
	Liaison Modem : Equipement ne redémarre plus	Deserver 90 HS	Arrêt définitif à confirmer avec le COMUTI	RB
	EXTRA ne redémarre pas automatiquement		Redémarrage manuel. Voir possibilité de redémarrage automatique	RB
Monitoring				
Serveurs Internes				
	Robot backup non détecté	Robot allumé après le démarrage du cluster (power off sur les alims du robot)	Suivre la procédure indiqué dans le FAQ	MJ
	www.lal.in2p3.fr inaccessible de l'extérieur	Mauvais redémarrage du CISCO	Redémarrage du CISCO	GP
	pcserv n'a pas redémarré	Raison non identifiée	Power off/on, monitoring de pcserv sous nagboard à vérifier	RB
	sbdd : echec du backup	Client Networker désinstallé au reboot par Quattor (installé manuellement)	Mise à jour de la config Quattor (grif:r13524)	MJ
	sdbb : non redémarrage des bases ZODB	Pas de script de démarrage dans / etc/init.d	Script positionné pour auto start, synchro à faire avec les pubs (de toute façon le downtime maxi est de 30 mn car les pub redémarrent leurs démons httpd et se resynchronisent dans ce délai, template sur Quattorsrv à updater	EM/GP
	pub2 : Démarrage de httpd au lieu de httpd-lal (mauvaise page d'acceuil)	Problème de config Quattor ?		EM
	Synchronisation macserv/ macserv2	macserv2 démarre trop vite, avant l'ouveture des bases par macserv	Redémarrer macserv2 après le démarrage de macserv. A documenter	BL
	Baie dell1 : Environ 7 srv sans pile BIOS	Pile trop ancienne	Vu avec Gerard pour changer les piles à la prochaine coupure si possible	GD/GP
Serveurs Grille				
	grid11 : impossible de se logger via gsissh	Locked accounts (poblème de configuration Quattor)	Mise à jour du profile quattor pour unlock les accounts (r13610)	GP

2008 Quatt... Doc/F... 10 S... 8 Mi... wms 3... dpm_... 2 Mi... Micro... Mulbe... FR 00:05

... Trac : Example

Diff r20374:20383 for trunk/cfg/grid - Tier 2 LCG Ile-de-France - Trac - Opera

Fichier Editer Afficher Signets Widgets Flux RSS Discussion Outils Aide

https://trac.lal.in2p3.fr/LCGTier2/changeset?new=trunk%2Fcfg%2Fgrid%4020383&old=trunk%2Fcfg%2Fgrid%4020374

logged in as /C=FR/O=CNRS/OU=UMR8607/CN=Michel Jouvin/emailAddress=jouvin@lal.in2p3.fr | Logout | Settings | Help/Guide | About Trac

Wiki | Timeline | Roadmap | Browse Source | View Tickets | New Ticket | Search | Tags | Admin | Discussion | Blog

Reverse Diff

Changes in trunk/cfg/grid [20374:20383]

Files:
- trunk/cfg/grid/glite-3.0.0/common/lcg/env.tpl (1 diff)
- trunk/cfg/grid/glite-3.1/glite/wn/rpms/i386/config.tpl (1 diff)
- trunk/cfg/grid/glite-3.1/common/lcg/env.tpl (1 diff)

Unmodified Added Removed Modified Copied Moved

View differences: inline

Show 2 lines around each change

Ignore:
- Blank lines
- Case changes
- White space changes

Update

trunk/cfg/grid/glite-3.0.0/common/lcg/env.tpl

r17087	r20383	
24	24	
25		Titre: Show new version of trunk/cfg/grid/glite-3.0.0/common/lcg/env.tpl
26		Adresse: https://trac.lal.in2p3.fr/LCGTier2/browser/trunk/cfg/grid/glite-3.0.0/common/lcg/env.tpl?rev=20383
27		"/software/components/profile/env/SRM_PATH" = DCACHE_LOCATION+"/srm";
	26	# Should not be defined anymore, causing a warning when defined.
	27	#variable DCACHE_LOCATION = INSTALL_ROOT+"/d-cache/";
	28	#"/software/components/profile/env/SRM_PATH" = DCACHE_LOCATION+"/srm";
28	29	
29	30	# Locations of "optional" machine types.

trunk/cfg/grid/glite-3.1/glite/wn/rpms/i386/config.tpl

r19928	r20383	
169	169	};
170	170	
	171	# Required by update 22
	172	'/software/packages'= if (GLITE_UPDATE_VERSION >= '22') {
	173	pkg_del('dcache-client','','i586');
	174	pkg_repl('dcache-client','1.8.0-0','noarch');
	175	pkg_repl('dcache-dcap','1.8.0-4','i586');

Other Collaborative Tools

- Every day communication is crucial
  - Peer-to-peer communication, not support with low-level sysadmins and experts
  - Allow irregular participation: nobody is 100% available
  - A place where new comers or people with involved for a small percentage can increase their expertise
- Currently based on an email list: doesn't scale
  - Some days with 100 messages..
- Looking at more agile tools
  - Daily meeting à la WLCG : probably too heavy and difficult to organize with enough participation
  - IRC channel
- Video-conference for participation to French meetings (LCG T2/T3 coordination, EGEE SA1-FR...)
  - GRIF sharing its expertise with other sites, eg. Storage (space tokens, ACLs, DPM support)

Quattor…

- Tool allowing to manage complex configuration in O(10) to O(1000) machines
  - Initial installation and every day (re)-configuration from same configuration information
  - Reproducible installations
  - Efficient, factorized configuration description
  - Client available for all platforms (pure Java), including Eclipse
  - Whole site can be administered from everywhere with only a SVN (https) connection
- Grid MW requires ability to quickly deploy frequent updates
  - OS, MW, CAs…
  - Avoid duplication of effort : manpower limited, site consistency
  - Requires possibility of staged deployment : cannot take the risk of breaking 5 or 6 CEs or SEs in a few seconds…

… Quattor

- GRIF specific need : manage a large number of machines spread over different location
  - GRIF currently 600+ machines in 6 sites
- Main goal: avoid duplication of effort, allow everybody to participate
  - 1 person can take in charge upgrading whole GRIF without a significant load/complexity increase
    - Deployment time of a configuration change : ~5mn
    - Possibility of staged deployment: 1 machine/cluster first…
  - PAN compilation allows to predict effects of changes
  - Quattor used for managing non grid machine allowing integration between grid and non-grid sysadmins
    - Allow to involve more persons
- GRIF is one of the major contributor to QWG effort
  - Others can benefit from features added in standard templates
  - But GRIF is not the only one… (at least 5 or 6 sites, increasing)
    - E.g. : support for Xen-based VMs

Quattor Workflow

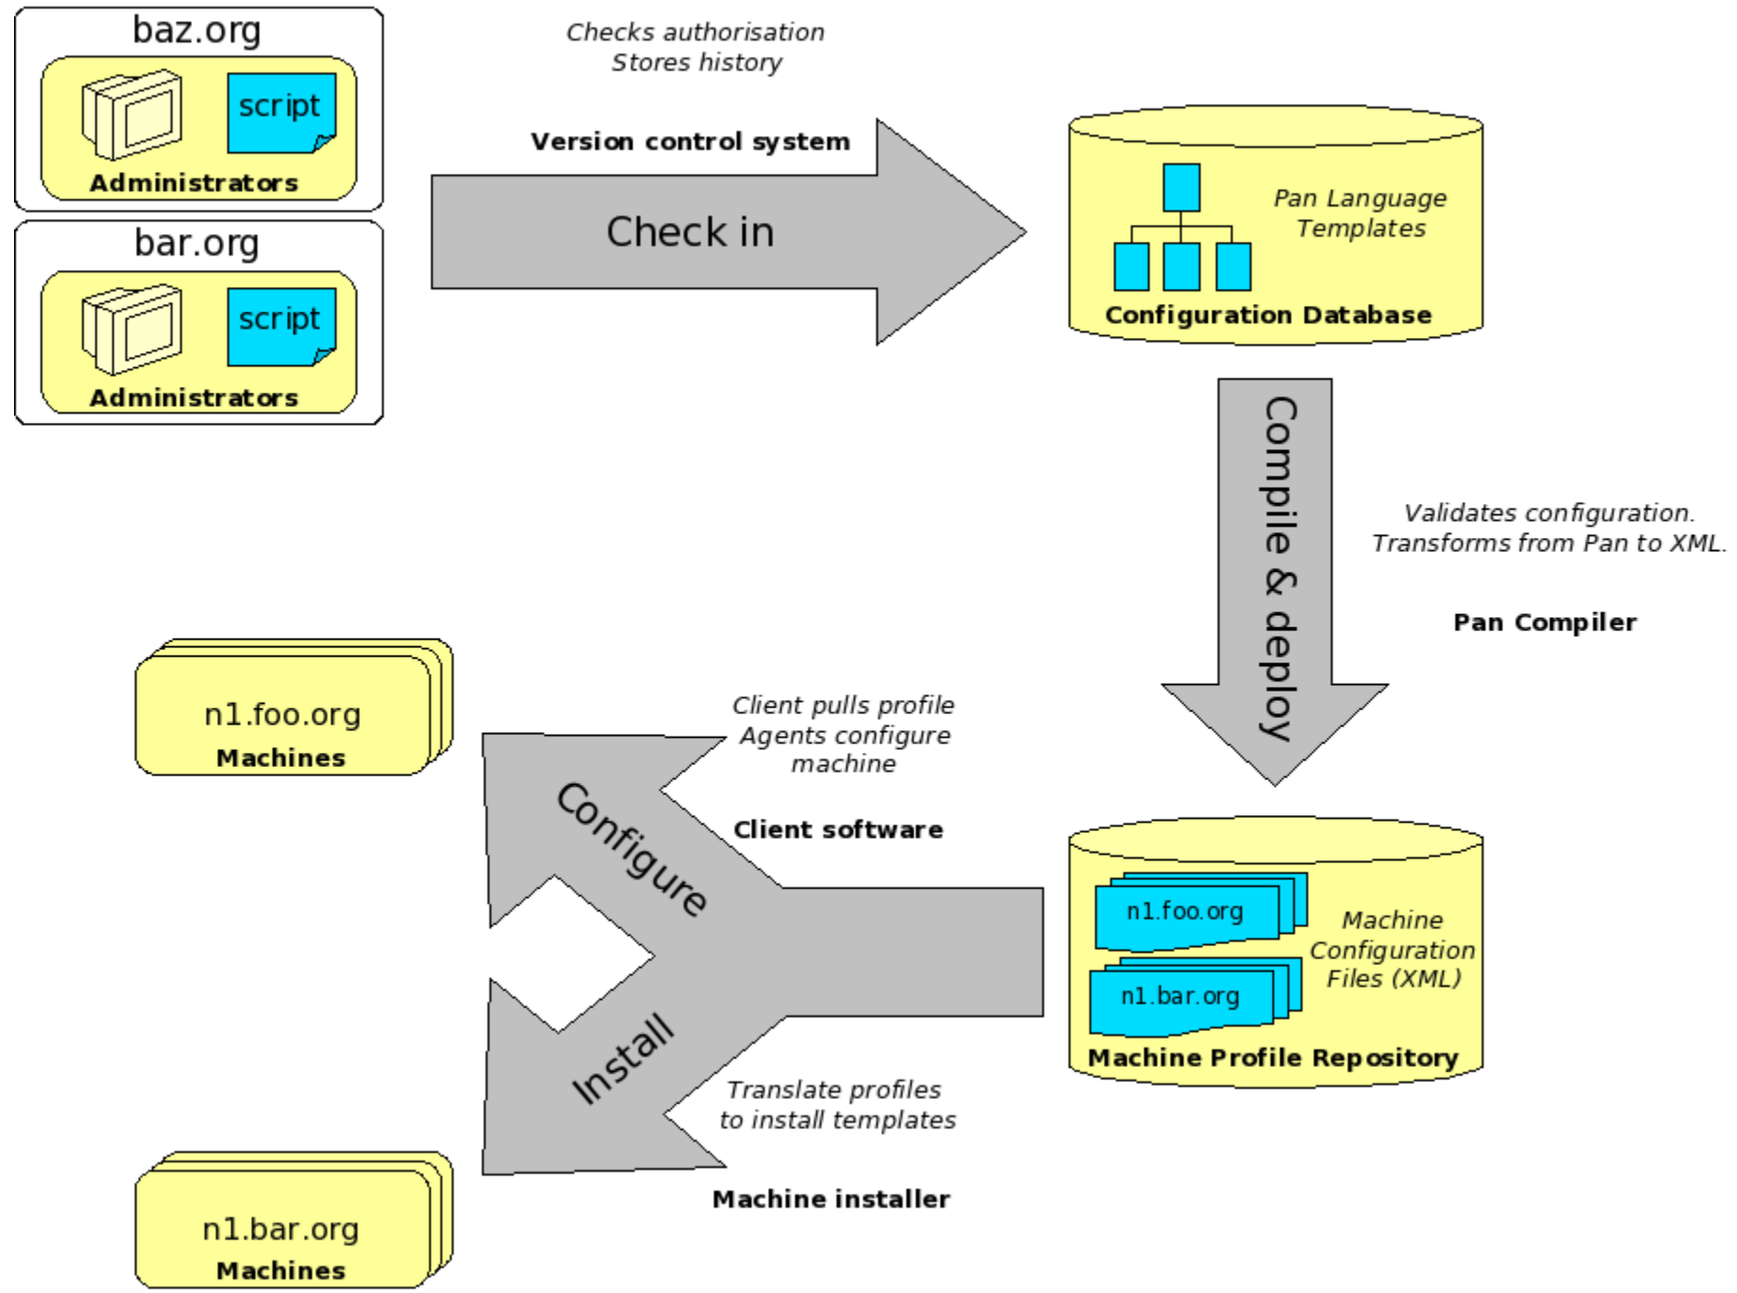

Relationships with Expts

- Great attention paid to establish links with experiments
  - 1-2 persons being the main contact for each experiment, spread over all the GRIF sites
- Experiments have difficulties with a distributed T2s running several SEs/CEs
  - 1 CE/SE per site means they look as different sites for data transfers
    - Had some problems with a FTS channel used by all Ses
      - Now fixed with a small change allowing 1 channel per « subsite » advertised in BDII
    - Moving to each SE being seen as close SE by all CEs
      - Specialize SE/VO to have 1 or 2 SE per VO with big chunks and no duplication
      - Doesn't fit with CMS catalog-less model
- Huge improvements in the last 6 months…
- Working with LHCb to handle some analysis tasks
  - Based on involvement of local LHCb physicists

Conclusions

- Succeeded to create a reliable distributed site managed by a large distributed team
  - Working with CMS for a definition of availability requiring all SEs opened to the VO
  - Having many people involved with enough expertise is a definitive advantage for a good coverage during 365 days
- Human dimension is the most critical
  - Everybody must be comfortable with his involvement
    - Nobody must feel that grid/GRIF leads to a less interesting work
      - More important when you use a tool like Quattor
      - Success is good for everybody and increases the “global dynamism”
- Tools are critical to compensate geographical dispersion
  - Facilitate “horizontal” communication between people
  - Allow traceability of configuration changes
  - Check/enforce configuration consistency across the whole site

Useful Links

- GRIF : http://grif.fr

- Trac :
  - Main site : http://trac.edgewall.org/
  - Plugins : http://trac-hacks.org/

- Quattor : https://trac.lal.in2p3.fr/LCGQWG