FRANCE GRILLES

Grid Operations : Status and perspectives

Gilles Mathieu, Hélène Cordier

France Grilles International Advisory Committee – April 2012

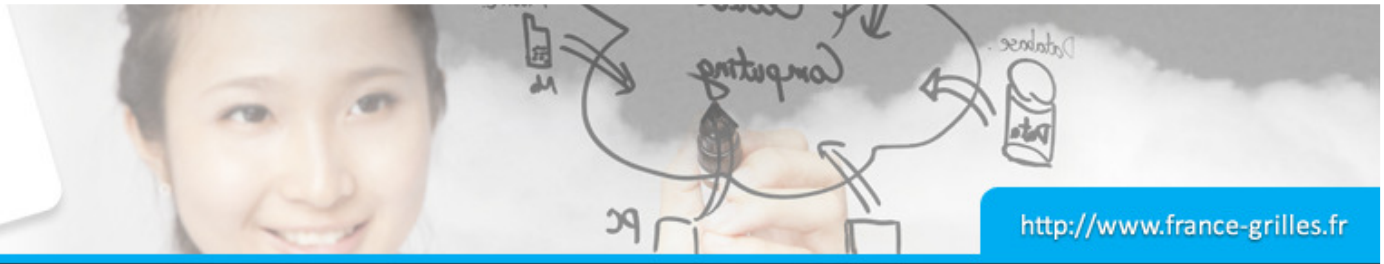

Outline

- Where are we ?
  - Some enlightening metrics
- Last year's highlights
  - Events and workshops
  - Catalogue of services
  - Internal tools and organization
  - Starting cloud initiatives
- Perspectives for the year to come
  - When Grid Operations meet Grid Users…
  - Getting more metrics: Methodology
  - Cloud perspectives
- Towards sustainability
  - Evaluating the cost and benefits of Grid Operations
  - Running national and global services beyond EGI-InSPIRE
  - What issues are we facing? How can we solve them?

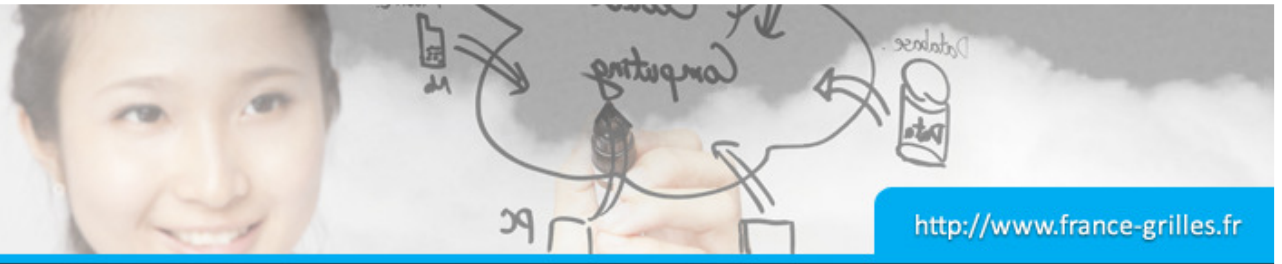

Outline

- Where are we ?
  - Some enlightening metrics

YOU ARE HERE

- Last year's highlights
  - Events and workshops
  - Catalogue of services
  - Internal tools and organization
  - Starting cloud initiatives
- Perspectives for the year to come
  - When Grid Operations meet Grid Users…
  - Getting more metrics: Methodology
  - Cloud perspectives
- Towards sustainability
  - Evaluating the cost and benefits of Grid Operations
  - Running national and global services beyond EGI-InSPIRE
  - What issues are we facing? How can we solve them?

Some enlightening metrics (1) – Available resources

- Available resources and their distribution
  - **18** production sites
  - **31,860** logical CPUs
  - **280,316** HEPSPEC06
  - **15.7 PB** disk storage
  - **22.5 PB** tape storage

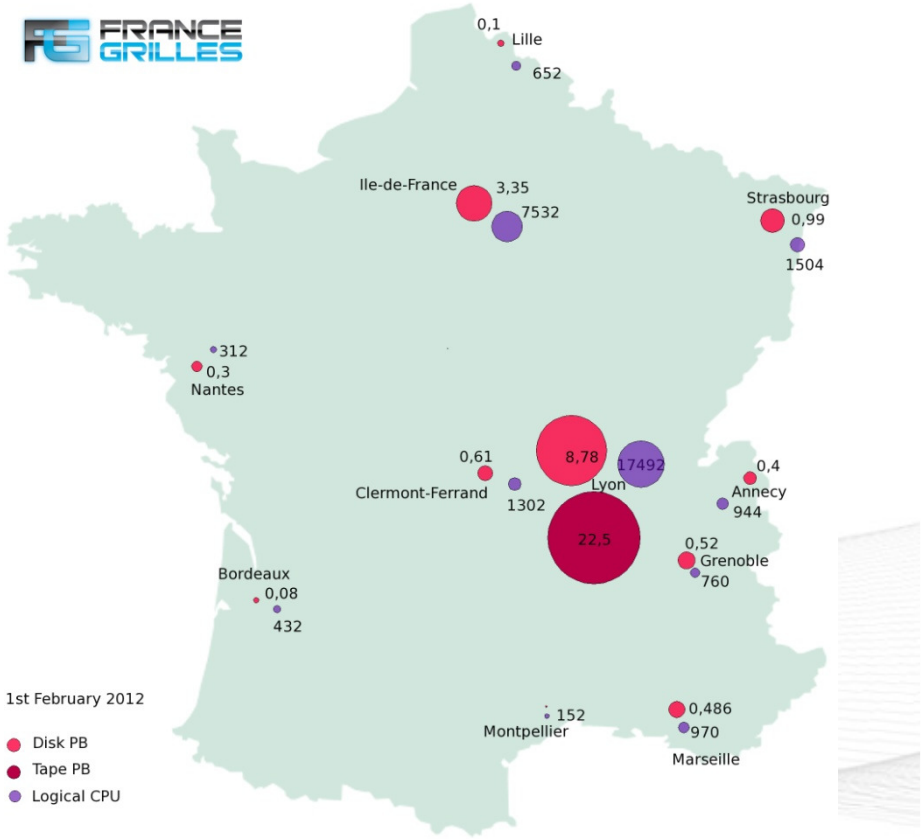

Some enlightening metrics (2) – Available resources

Logical CPUs over the last year

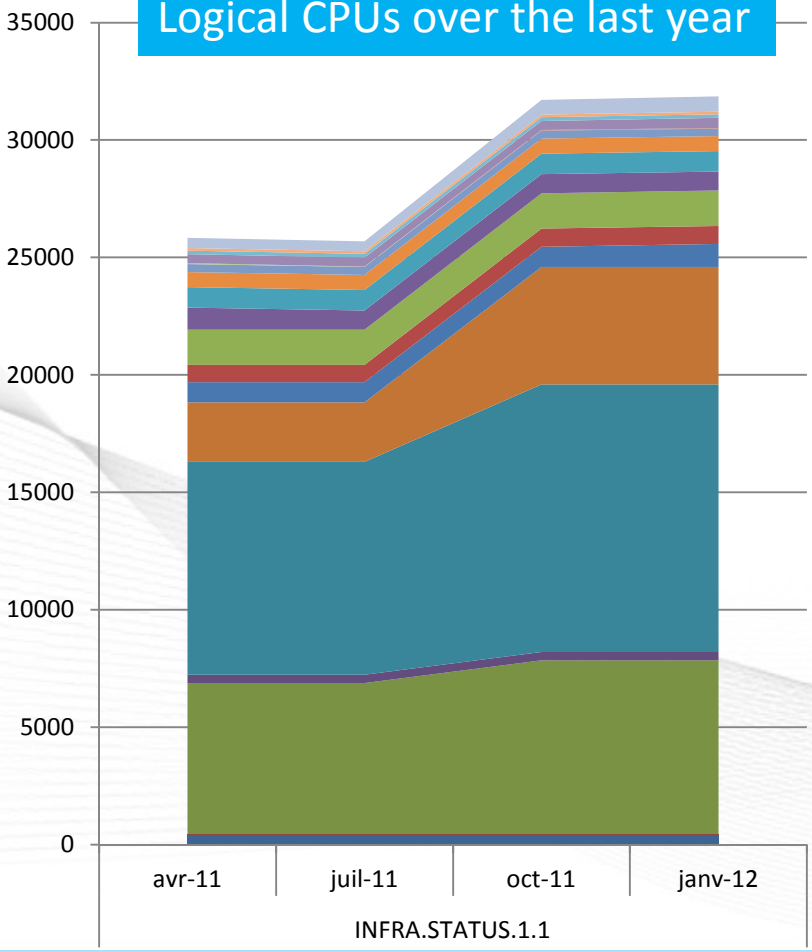

Disk storage (PB) over the last year

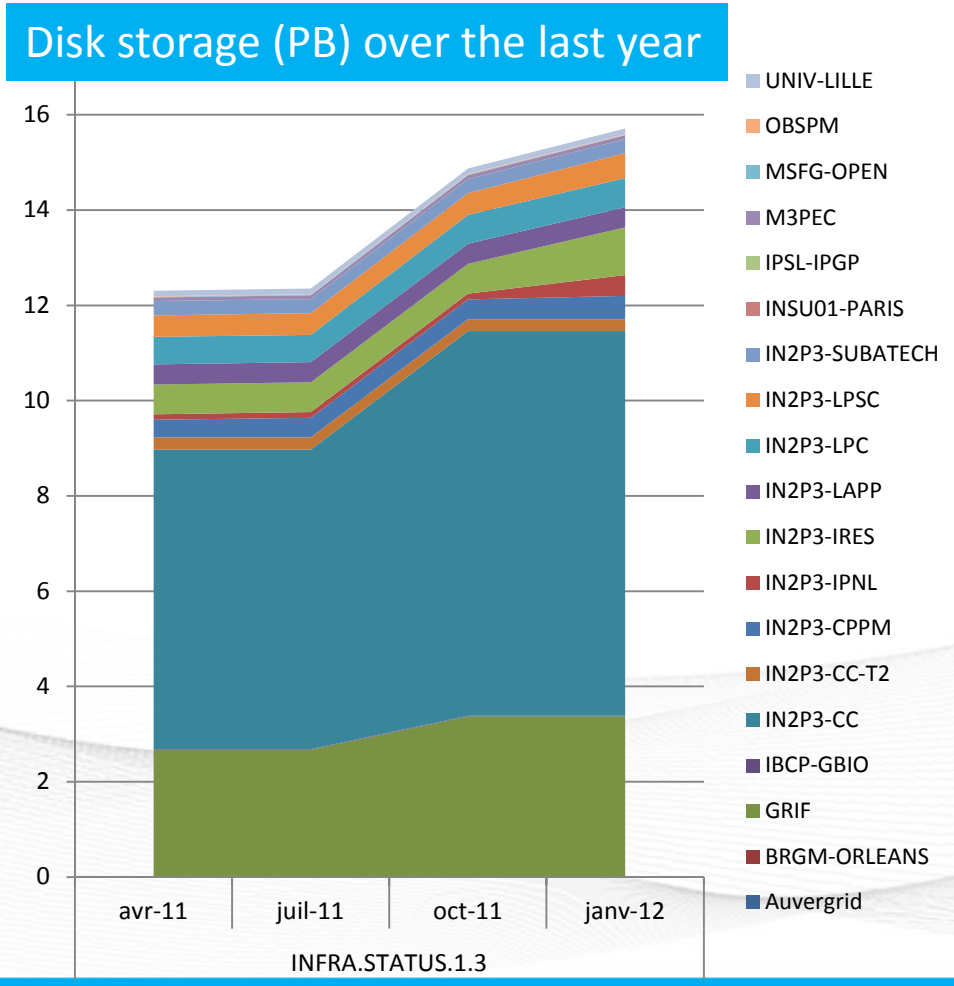

Some enlightening metrics (3) – Used resources

- Used resources

Normalized CPU time (HEPSPEC06.hour) per quarter since Q1

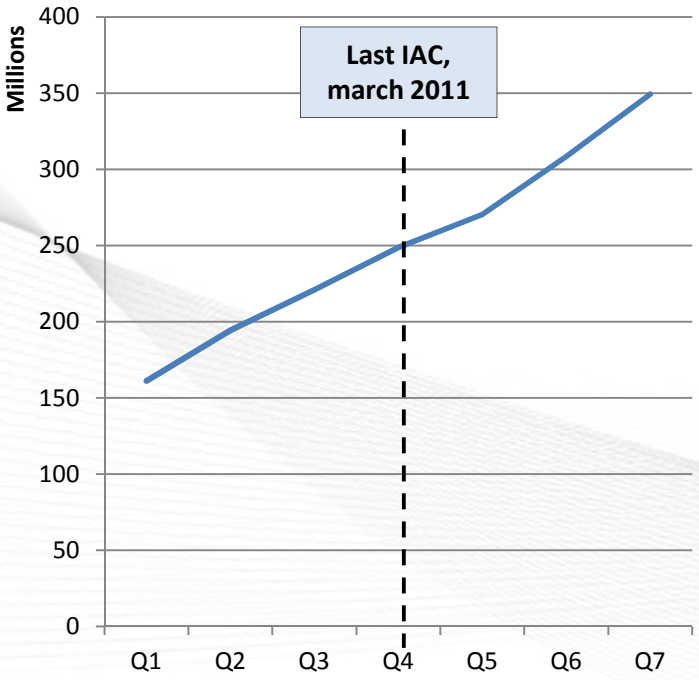

Normalized CPU time per site from Q1 to Q7

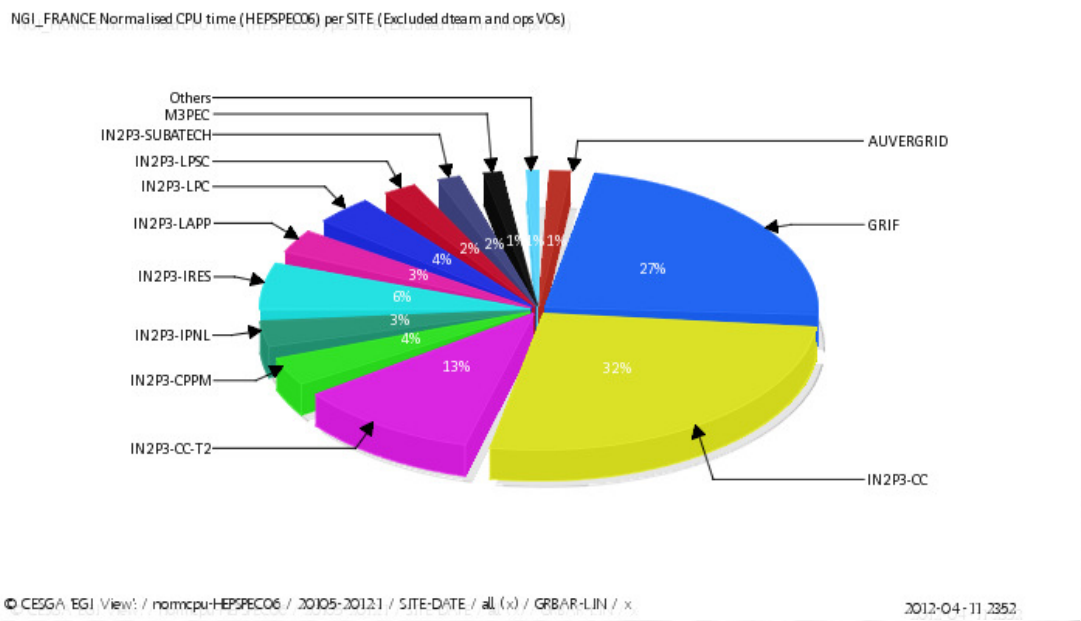

Some enlightening metrics (4) – French contribution to EGI

- Computing resources
  - 11.1% of available capacity
  - 14.4% of used capacity
    - 3rd contributor to EGI behind:
      - UK (16.6%)
      - Germany (15.4%)
- Storage resources
  - 11.8% of available disk space
    - 5th contributor to EGI behind
      - Germany (14.8%)
      - CERN (14.4%)
      - UK (12.4%)
      - Italy (12.0%)

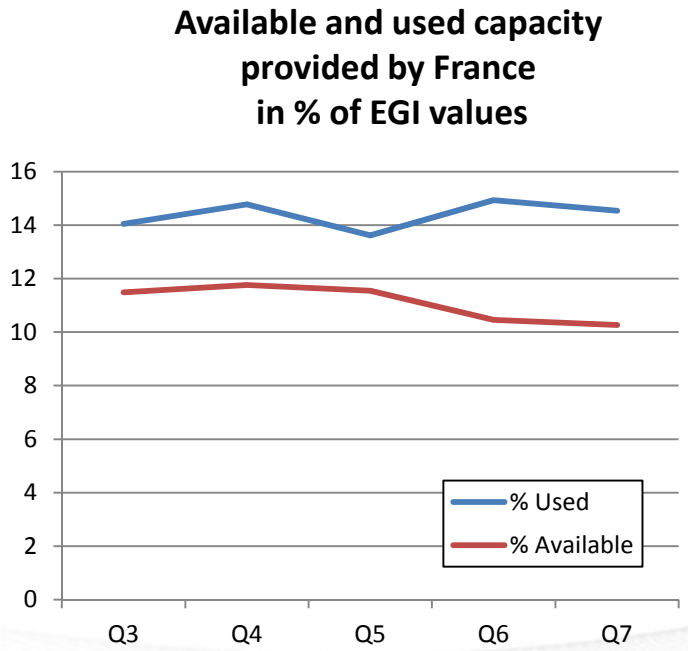

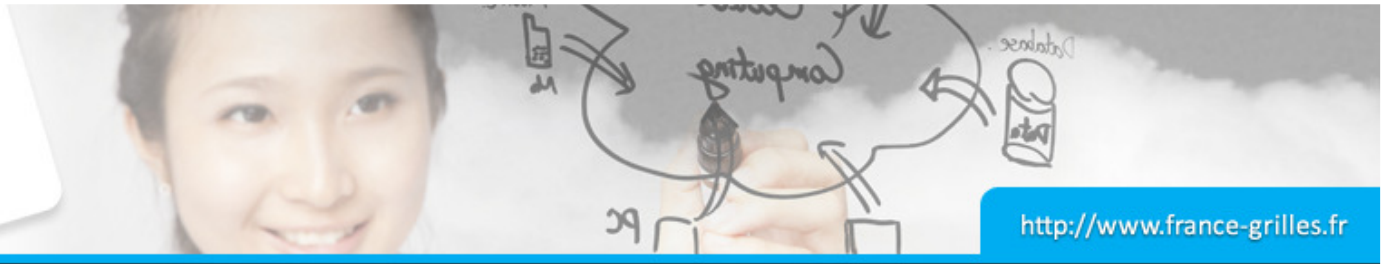

Some enlightening metrics (5) – Quality

- Quality of our infrastructures
  - **98.4%** availability
  - **99.0%** reliability
  - **100.0%** availability/reliability of our central core services (Top-BDII) since September 2011

Production infrastructure (average 04/2011 - 03/2012)

- Definitions of Availability and Reliability:
  - A = *(Up period)/(known period)*
  - R = *(Up period)/(known period – scheduled downtime)*
  - More on https://wiki.egi.eu/wiki/Availability_and_reliability_monthly_statistics

Some enlightening metrics (6) – Load and efficiency

- CPU efficiency
  - Sum (CPUtime)/Sum (Walltime)
  - **84%** average on last year

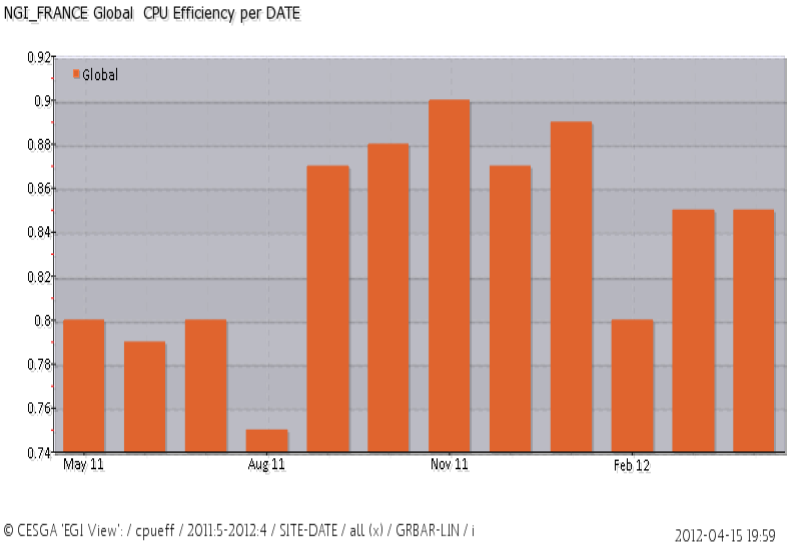

- Load on grid machines
  - Simplistic evaluation: (Used/Available)
  - Does that make sense? Still many questions to address
    - Ongoing work to determine more precisely how loaded is our infrastructure
    - Metric under definition

Some enlightening metrics (7) – Security

- Number of detected Critical Vulnerabilities (CVE)

EGI
As of 13/04: **378** CVE
Average: **270/day**

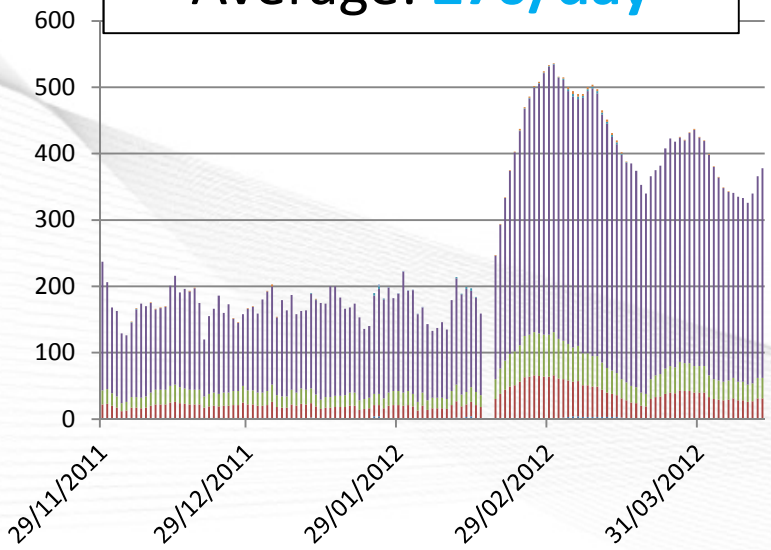

France
As of 13/04: **no** CVE
Average: **0.6/day**

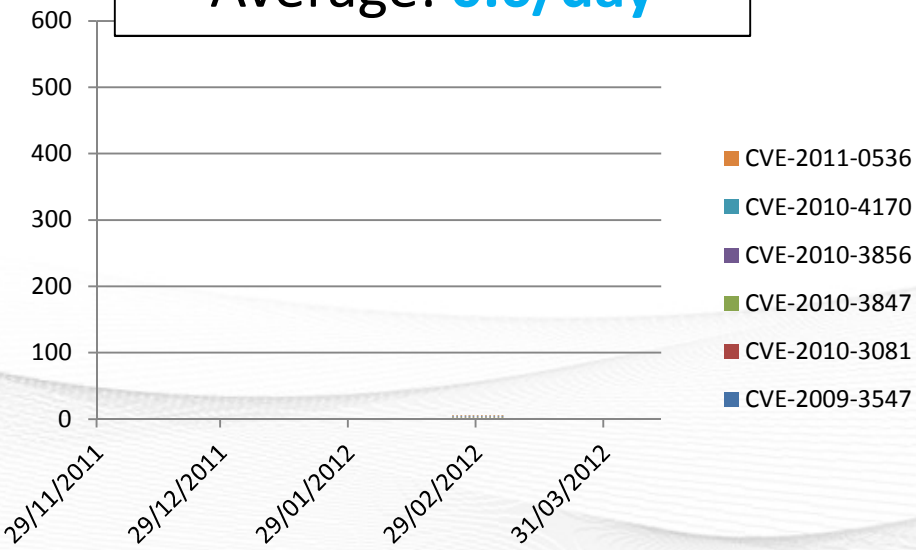

Source : Operations Portal/Security Dahsboard

Some enlightening metrics (8) – Operational support

- Tickets over the last year (Apr 2011 - March 2012)
  - **691** tickets assigned to NGI_FRANCE, **558** solved
  - Solution time: **12.5 days** average, **1.7 day** median

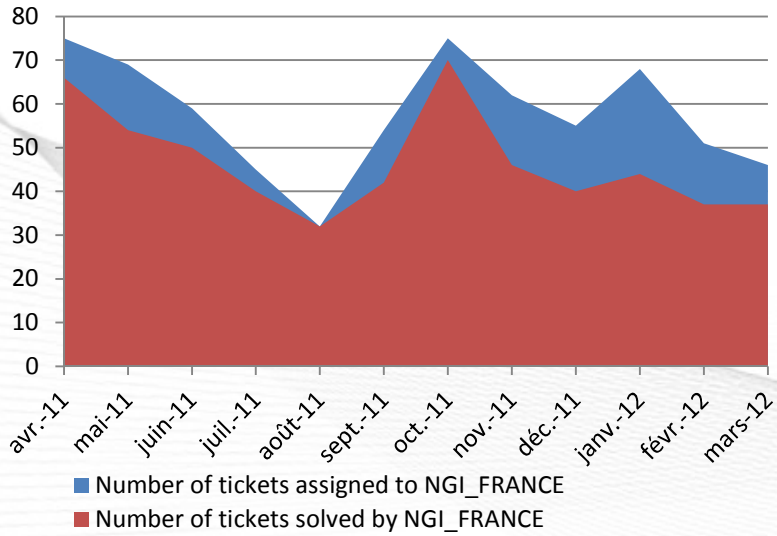

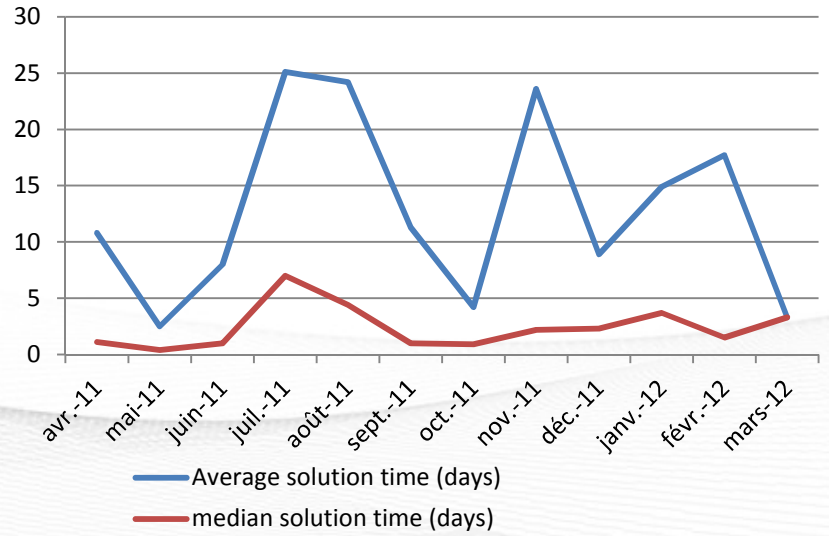

Source : GGUS / Global Grid User Support

Some enlightening metrics (9) – Effort and time

- Effort needed and spent on "central" tasks
  - Tasks at national/NGI level

	NEEDED EFFORT						CURRENT EFFORT	
	Crit.		Crit. + Import.		All			
	h/week	FTE	h/week	FTE	h/week	FTE	h/week	FTE
Technical coordination	45	1,18	71	1,87	88	2,32	41	1,08
Executive Technical committee	11	0,29	22	0,58	23	0,61	13	0,34
Accounting	19	0,5	23	0,61	27	0,71	18	0,47
Monitoring	23	0,61	35	0,92	36	0,95	17,5	0,46
Network	3	0,08	41,5	1,09	42,5	1,12	23	0,61
Certification Authority	36,5	0,96	52	1,37	53,5	1,41	37	0,97
Security	17	0,45	25	0,66	31	0,82	16	0,42
Deployment	2	0,05	11	0,29	41	1,08	38	1
Others	30	0,79	36	0,95	48	1,26	36	0,95
TOTAL	**186,5**	**4,91**	**316,5**	**8,34**	**390**	**10,28**	**239,5**	**6,3**



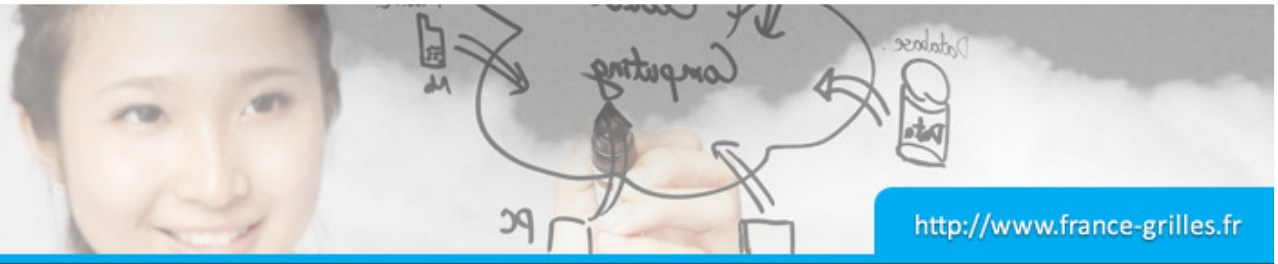

Outline

- Where are we ?
  - Some enlightening metrics
- Last year's highlights
  - Events and workshops
  - Catalogue of services
  - Internal tools and organization
  - Starting cloud initiatives
- Perspectives for the year to come
  - When Grid Operations meet Grid Users…
  - Getting more metrics: Methodology
  - Cloud perspectives
- Towards sustainability
  - Evaluating the cost and benefits of Grid Operations
  - Running national and global services beyond EGI-InSPIRE
  - What issues are we facing? How can we solve them?

YOU ARE HERE

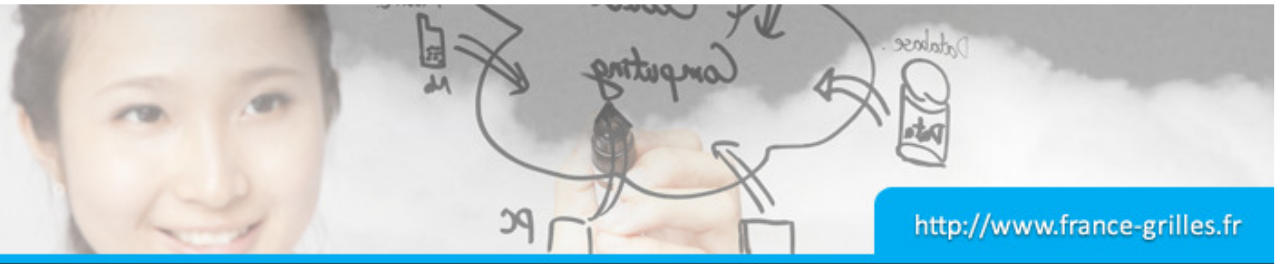

Events and workshops (1)

- 2nd Grid Operations Workshop – Oct. 20th 2011
  - Addressed topics
    - Placing operations within overall strategy
    - Internal organization of Grid Operations
    - France Grilles' service offer
  - A good occasion to get a different perspective
    - Presentations and discussions with UK and Polish NGIs
    - Broadening the scope by linking operations with the rest
  - Building a roadmap for the following 6 months

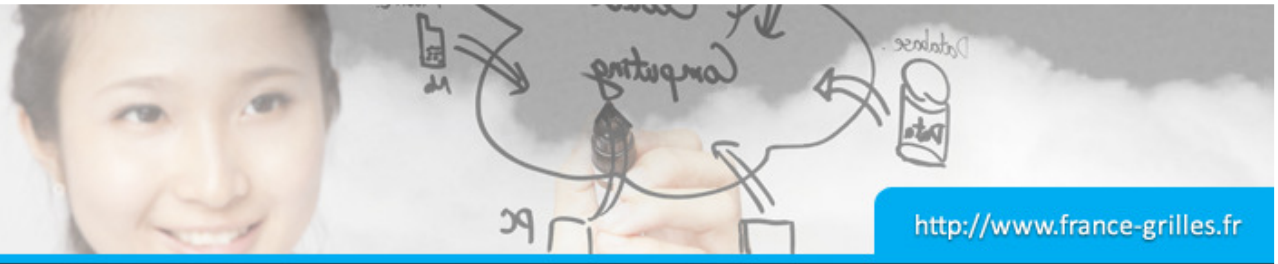

Events and workshops (2)

- France Grilles 2nd Cloud Day – Oct. 21st 2011
  - Addressed topics
    - Sharing and disseminating knowledge on cloud technologies
    - Assessing state-of-the-art in the area
    - Discussing issues and opportunities
  - Reached goals
    - Raising cloud awareness and knowledge
    - Paving the way for French cloud initiatives and wider implications

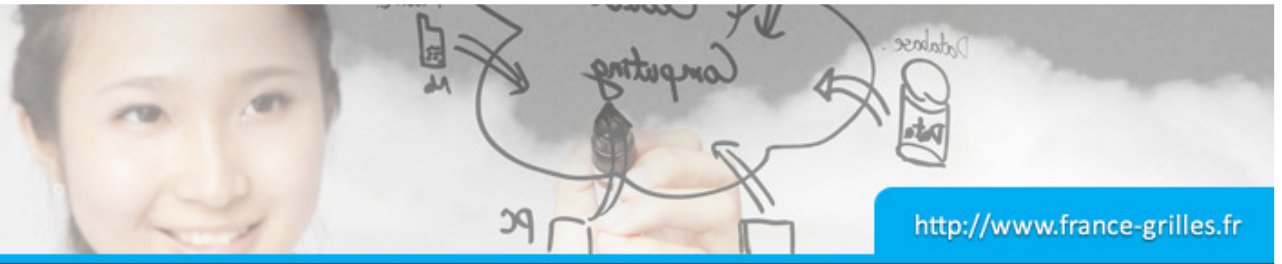

Events and workshops (3)

- Trainings for site administrators
  - ACGRID3 school, Vietnam, oct.2011 (53 participants)
  - Quattor administration, dec.2011 (16 participants)
  - DIRAC administration, jan.2012 (15 participants)

- EGI Technical Forum 2011
  - High implication from operational actors because the EGI-TF was in Lyon

Catalogue of services – Catch-all VO vo.france-grilles.fr

- Deployement and status
  - Deployment Launched April 2011
    - Going through questionnaires, feasibility study, interactions with regional VOs, AUP definition etc.
  - Technically in place since July 2011
  - 4 VO managers

- Resources and services
  - Now supported on 12 sites
  - Central services (LFC, VOMS) hosted at CC-IN2P3
  - Available through DIRAC framework (see next slide)

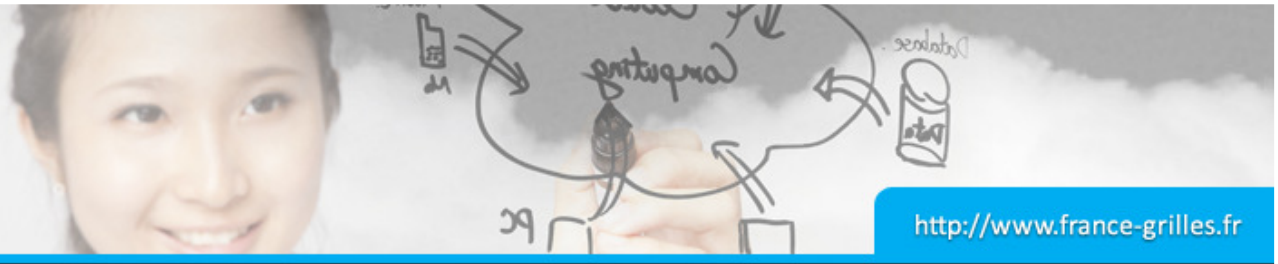

Catalogue of services – Providing new services (1) : DIRAC

- The DIRAC project
  - **D**istributed **I**nfrastructure with **R**emote **A**gent **C**ontrol
  - Initially developed for LHCb experiment
  - http://www.diracgrid.org

- A good opportunity for France Grilles
  - IN2P3 highly involved in development (CPPM, Marseille)
  - Seen as a great added value for our users

Catalogue of services – Providing new services (2) : DIRAC

- FG-DIRAC: France Grilles' national instance
  - A 3 phases project
    - Initialization (September 2011 – January 2012)
    - Preparation (January 2012 – April 2012)
    - Production (May 2012 onwards)
  - A multi-tier collaboration

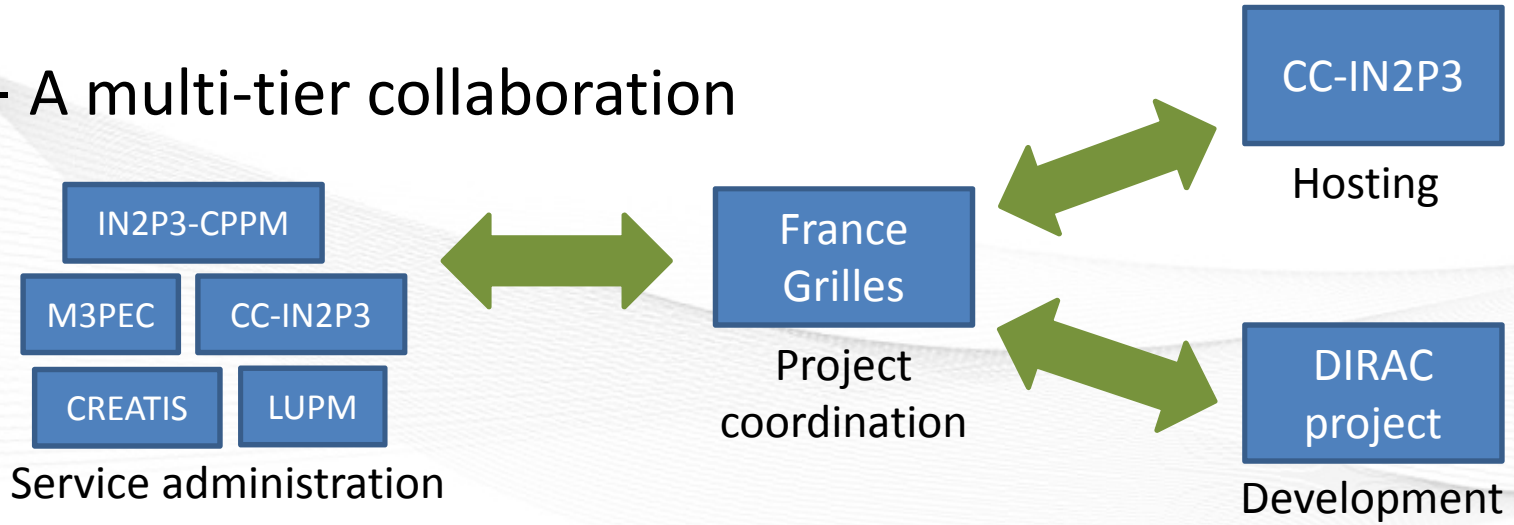

Catalogue of services – Providing new services (3) : DIRAC

- FG-DIRAC: current status
  - Ready to start production as scheduled
    - Platform installed and operational
    - MoUs underway (for hosting and administration)
    - Preliminary tests successful
    - Pilot communities ready to use the service
  - Ready to have it used and to communicate about it
    - Advanced dissemination at next operations workshop
    - Admin training in June

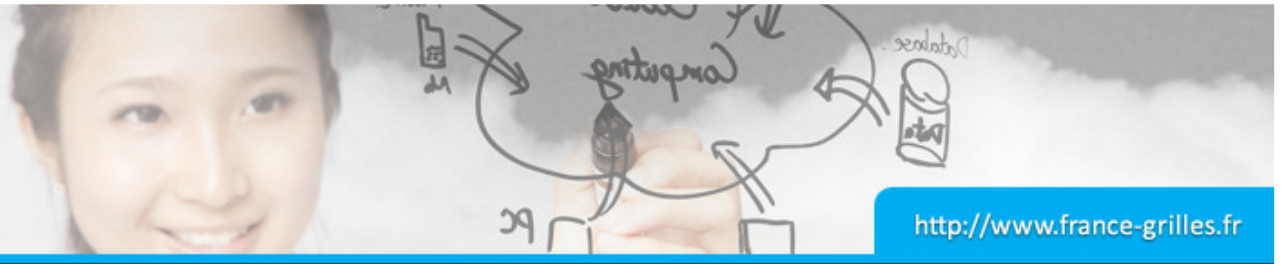

Catalogue of services – Providing new services (4)

- iRODS storage system
  - Plans to open the service to FG community
  - Discussions planed at next operations workshop
  - Admin training scheduled in June

- Tools Catalogue
  - Software used and/or developed within a Grid context
  - https://www.projet-plume.org/grille

Internal tools and organization (1)

- Operational Model : formalizing missions and tasks

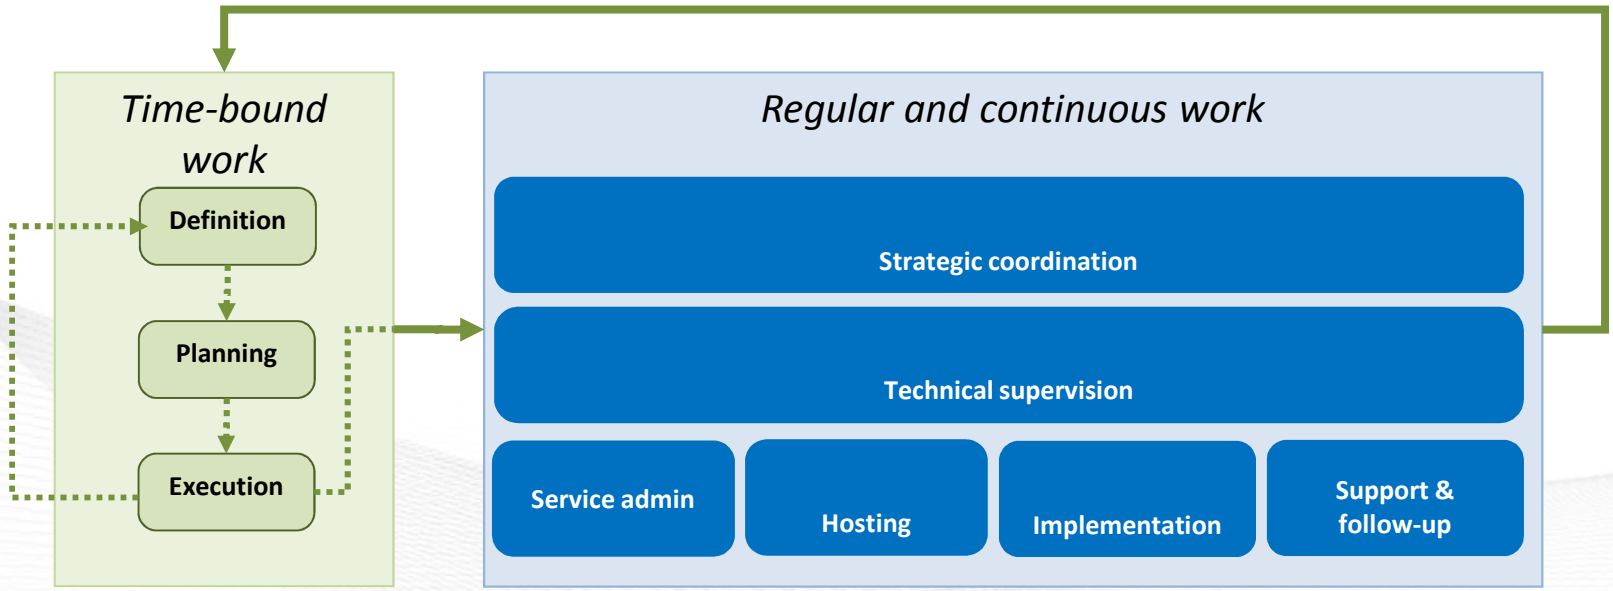

Internal tools and organization (2)

- Formalizing the definition of our infrastructures

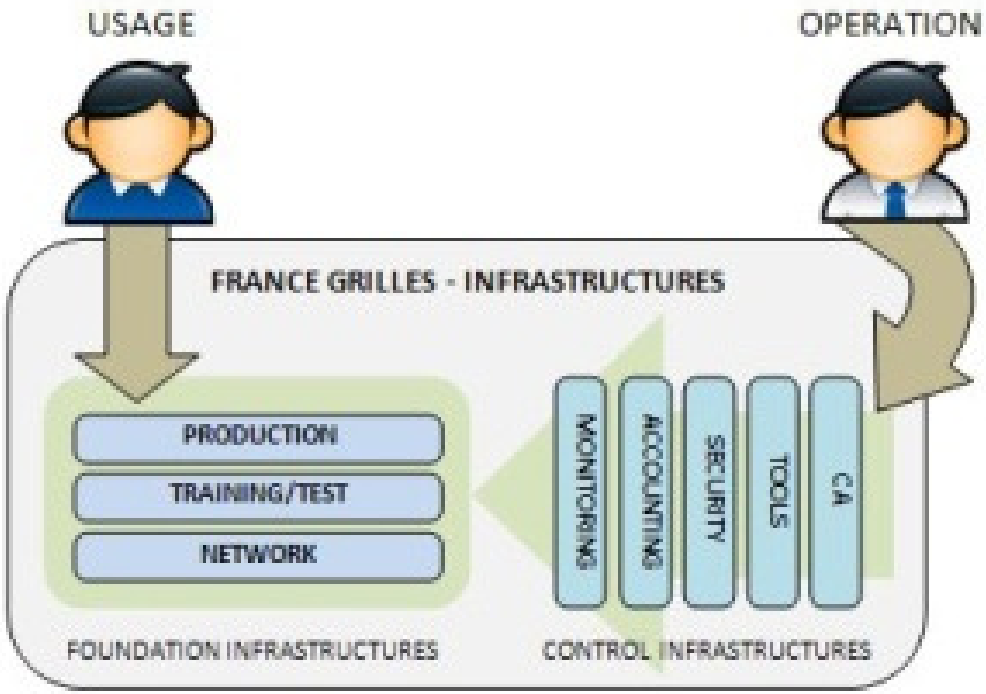

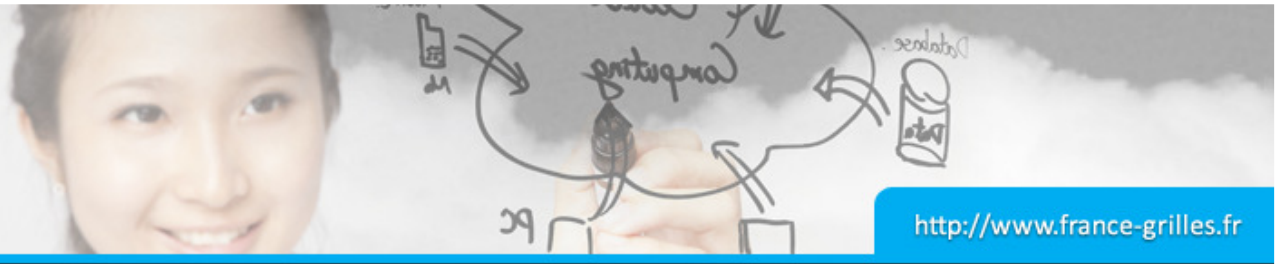

Internal tools and organization (3)

- Setup of an Executive Technical Committee
  - Experts + representatives from site and user communities
  - Discussing and validating strategic orientations
- Expert groups
  - Accounting, Monitoring, Security, Deployment, Network, Certification Authority
- Collaborative tools
  - Wiki, task tracker, RPM repository

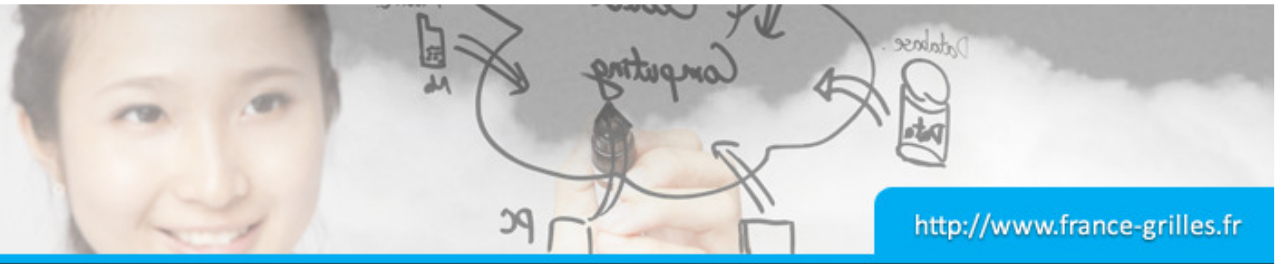

Cloud initiatives (1) – Academic Cloud at CC-IN2P3

- Who/what
  - Initiated in Nov. 2011 by Mattieu Puel
  - France-Grilles directly sponsored : 384 cores (16*24 cores)+ 1,5 TB (16*96 GB) +32 TB
- State of the art
  - Currently evaluating Openstack/Opennebula/Vmware
  - Evaluation criteria:
    - General considerations (wide adoption, compliance with standards, scalability, price)
    - functionalities (cloud management, monitoring, accounting, local storage)

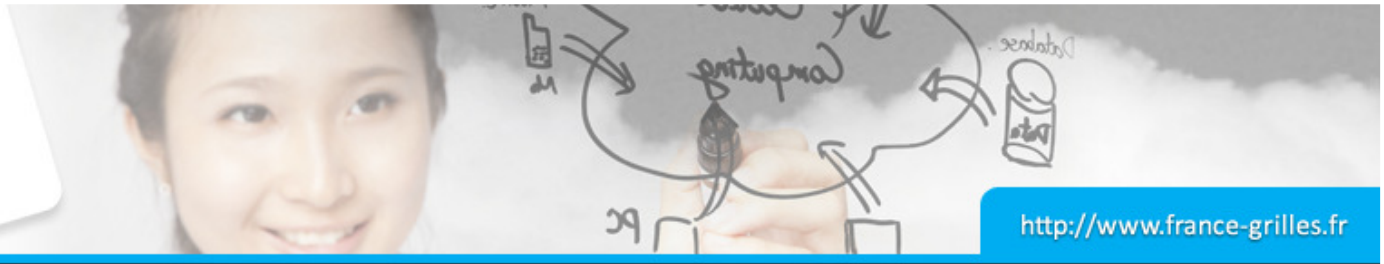

Cloud initiatives (2) – Academic Cloud at CC-IN2P3

- Steps
  - Offer IAAS resources through generic interfaces (EC2/OCCI)
  - Integrate national/international academic federations
- Use-case to be opened to :
  - EGI Cloud Federation Task Force testbed
  - DIRAC usage of cloud resources
  - Individual early users (biomed)
  - Webimatics (neurobiology)
  - Later: eTRIKS project (drug discovery)
- Participation to EGI Cloud Federation TF
  - WP4 (development of the Openstack driver for federated accounting)
  - WP5 (VM monitoring)

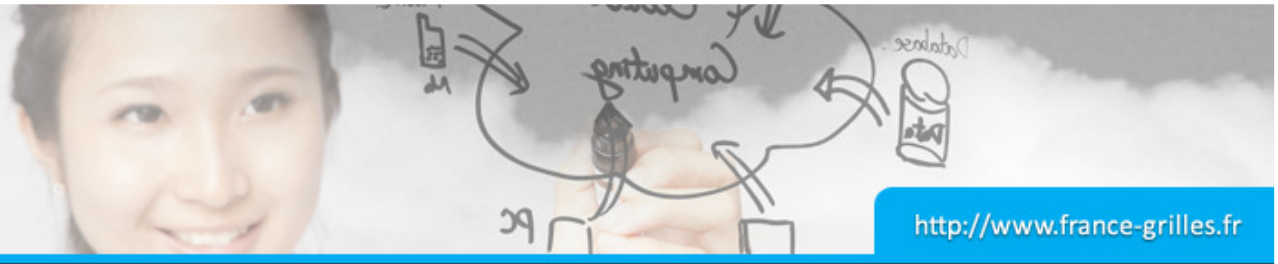

Cloud initiatives (3) – Hybrid Cloud at cloudmip.univ-tlse3.fr

- Who/what
  - Managed by F. Thiebolt & J.M. Pierson
  - France-Grilles directly sponsored : 32 blades (8 cores, 32Go RAM, 2 x SAS 15ktpm) – March 2012

- State of the art :
  - Scientific Linux 6.2 + KVM + OpenNebula, or
  - OpenStack (IaaS) + GlusterFS (RedHat DFS) + CompatibleOne (SaaS)
  - Alternative : RHEV3.0 + GlusterFS

Cloud initiatives (4) – Hybrid Cloud at cloudmip.univ-tlse3.fr

- Interests /issues
  - Green IT in datacenter
  - Cloud & inter-cloud distributed storage
- Use-case to be opened to :
  - Local academics
  - Users of grid5K
  - Some private companies

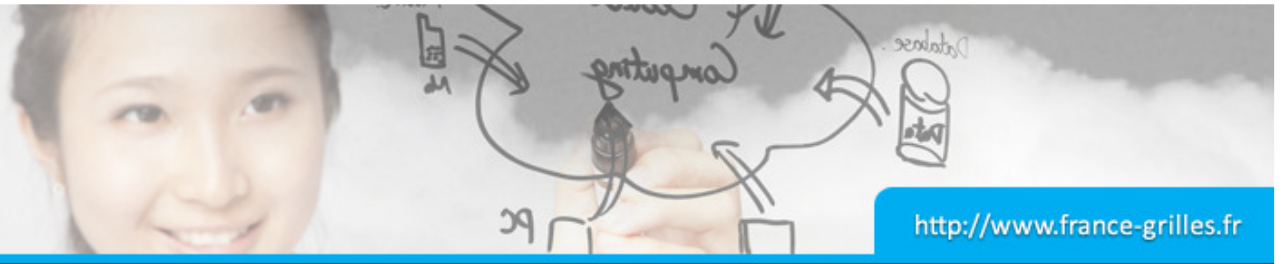

Outline

- Where are we ?
  - Some enlightening metrics
- Last year's highlights
  - Events and workshops
  - Catalogue of services
  - Internal tools and organization
  - Starting cloud initiatives
- Perspectives for the year to come
  - When Grid Operations meet Grid Users…
  - Getting more metrics: Methodology
  - Cloud perspectives
- Towards sustainability
  - Evaluating the cost and benefits of Grid Operations
  - Running national and global services beyond EGI-InSPIRE
  - What issues are we facing? How can we solve them?

YOU ARE HERE

Does that make sense to have plans ?

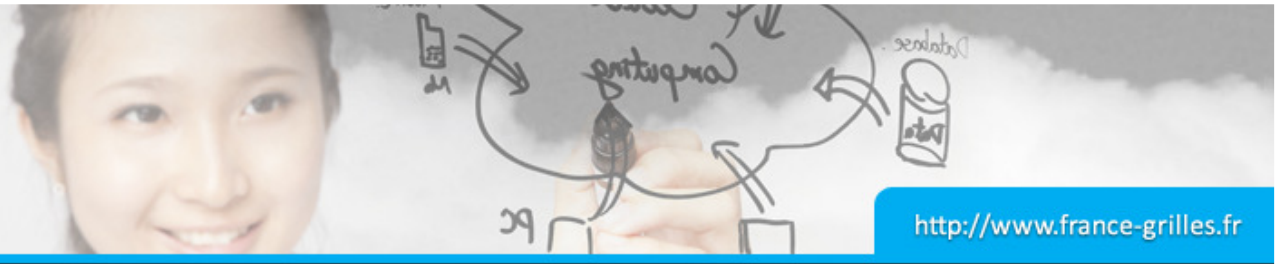

What next? (1)

- Keep on going…
  - Maintain the Quality of Service we already deliver
  - Bring FG-DIRAC to a full production service
  - Continue towards adding iRODS to our catalogue

- 3rd FG Operations workshop – Spring 2012
  - May 10th in Lille
  - Advanced dissemination on DIRAC and iRODS
  - Engaging operations with user communities

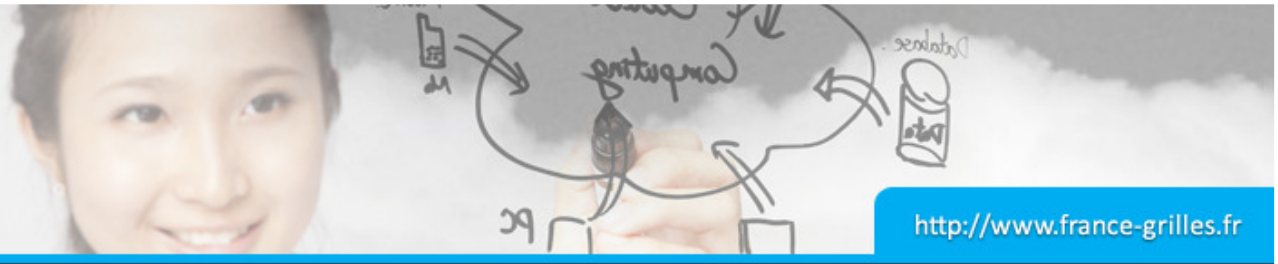

What next? (2)

- Strengthen relations between operations and:
  - User communities
    - Encourage collaborative support
    - Provide user training on new services
    - Federate help on application porting
  - Grid sites
    - Formalize relationships, establish MoUs
    - Start thinking in an "SLA" way

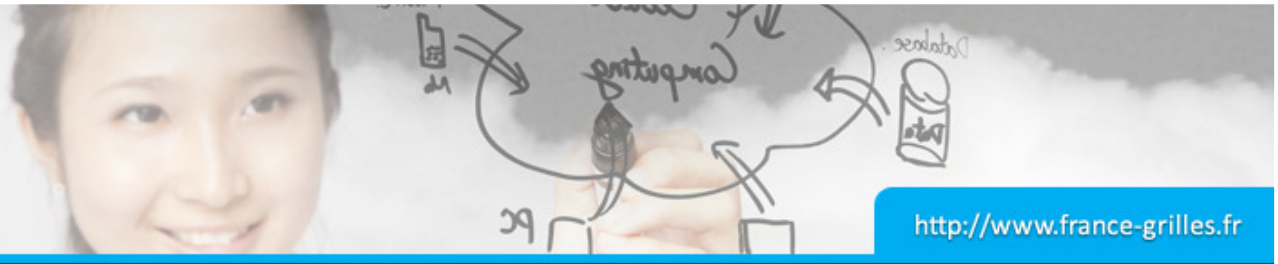

Thou shalt produce metrics (1)

- Metrics definition methodology
  - Work bottom-up
    - from *"what do we want to know"*…
    - … to *"what metric do we need for that"*
  - Define areas by origin : who produces the metric?
    - Operations, user communities, training, communication, budget/finances, collaborations/projects
  - Standardize description for each metric
    - Metric definition model

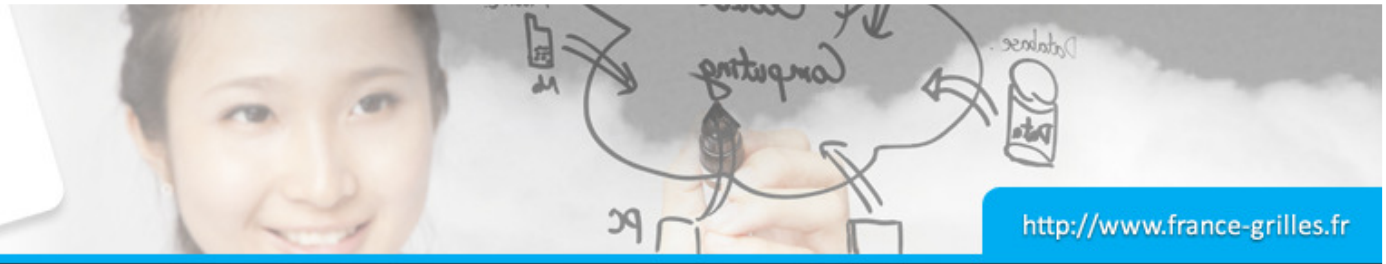

Thou shalt produce metrics (2)

- Metrics definition model

Category	(e.g. Infrastructure, users, training, etc.)
Parent indicator group	(e.g. "available resources", "Operational support activities", etc.)
Indicator ID and name	(e.g. "INFRA-1.1 - Number of available logical CPU"...)
Scope	(e.g. national, EGI, ...)
Goal	Why do we want to measure that
Description	Detailed description of what is measured
Calculation method	How are number obtained, especially if they derive from other metrics
collection/publication frequency	(e.g. "Monthly retrieval, published when retrieved")
Sources and tools	Where do data come from, how they are collected
Responsible person	Person that is responsible for collecting data
target value	Could be a minimum value, an ideal value, a general tendency…
Report model	(e.g. results table, graph, raw numbers, …)
Status	(not started, in progress, defined, implemented...)
Collection Starting date	(e.g. 01/01/2010, PM1, Q3, etc.)

Thou shalt produce metrics (3)

- Status
  - Work in progress on operations metrics
  - "dashboard" on operations wiki
  - Most of the data already there

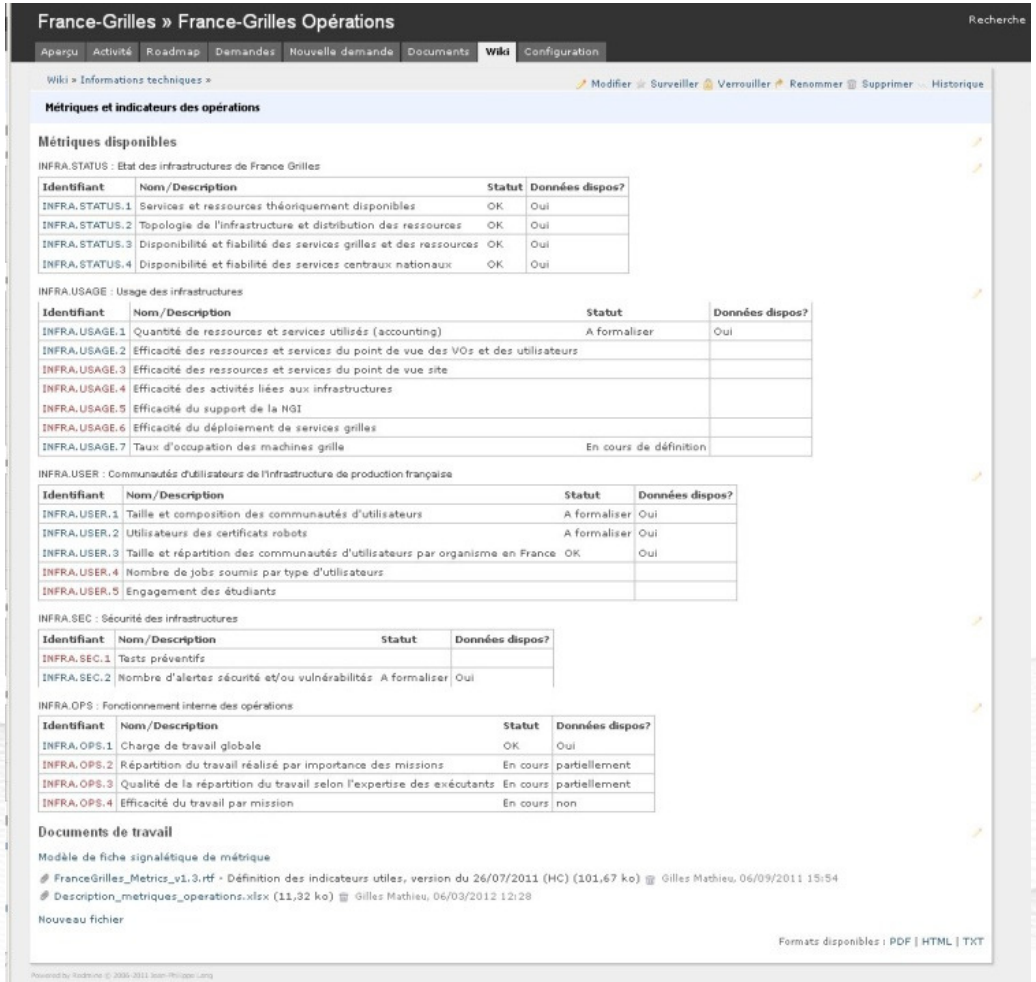

France-Grilles » France-Grilles Opérations

Recherche

Aperçu | Activité | Roadmap | Demandes | Nouvelle demande | Documents | Wiki | Configuration

Wiki » Informations techniques »

Modifier | Surveiller | Verrouiller | Renommer | Supprimer | Historique

Métriques et indicateurs des opérations

Métriques disponibles

INFRA.STATUS : Etat des infrastructures de France Grilles

Identifiant	Nom/Description	Statut	Données dispos?
INFRA.STATUS.1	Services et ressources théoriquement disponibles	OK	Oui
INFRA.STATUS.2	Topologie de l'infrastructure et distribution des ressources	OK	Oui
INFRA.STATUS.3	Disponibilité et fiabilité des services grilles et des ressources	OK	Oui
INFRA.STATUS.4	Disponibilité et fiabilité des services centraux nationaux	OK	Oui

INFRA.USAGE : Usage des infrastructures

Identifiant	Nom/Description	Statut	Données dispos?
INFRA.USAGE.1	Quantité de ressources et services utilisés (accounting)	A formaliser	Oui
INFRA.USAGE.2	Efficacité des ressources et services du point de vue des VOs et des utilisateurs		
INFRA.USAGE.3	Efficacité des ressources et services du point de vue site		
INFRA.USAGE.4	Efficacité des activités liées aux infrastructures		
INFRA.USAGE.5	Efficacité du support de la NGI		
INFRA.USAGE.6	Efficacité du déploiement de services grilles		
INFRA.USAGE.7	Taux d'occupation des machines grille	En cours de définition	

INFRA.USER : Communautés d'utilisateurs de l'infrastructure de production française

Identifiant	Nom/Description	Statut	Données dispos?
INFRA.USER.1	Taille et composition des communautés d'utilisateurs	A formaliser	Oui
INFRA.USER.2	Utilisateurs des certificats robots	A formaliser	Oui
INFRA.USER.3	Taille et répartition des communautés d'utilisateurs par organisme en France	OK	Oui
INFRA.USER.4	Nombre de jobs soumis par type d'utilisateurs		
INFRA.USER.5	Engagement des étudiants		

INFRA.SEC : Sécurité des infrastructures

Identifiant	Nom/Description	Statut	Données dispos?
INFRA.SEC.1	Tests préventifs		
INFRA.SEC.2	Nombre d'alertes sécurité et/ou vulnérabilités	A formaliser	Oui

INFRA.OPS : Fonctionnement interne des opérations

Identifiant	Nom/Description	Statut	Données dispos?
INFRA.OPS.1	Charge de travail globale	OK	Oui
INFRA.OPS.2	Répartition du travail réalisé par importance des missions	En cours	partiellement
INFRA.OPS.3	Qualité de la répartition du travail selon l'expertise des exécutants	En cours	partiellement
INFRA.OPS.4	Efficacité du travail par mission	En cours	non

Documents de travail

Modèle de fiche signalétique de métrique

FranceGrilles_Metrics_v1.3.rtf - Définition des indicateurs utiles, version du 26/07/2011 (HC) (101,67 ko) Gilles Mathieu, 06/09/2011 15:54

Description_metriques_operations.xlsx (11,32 ko) Gilles Mathieu, 06/03/2012 12:28

Nouveau fichier

Formats disponibles : PDF | HTML | TXT

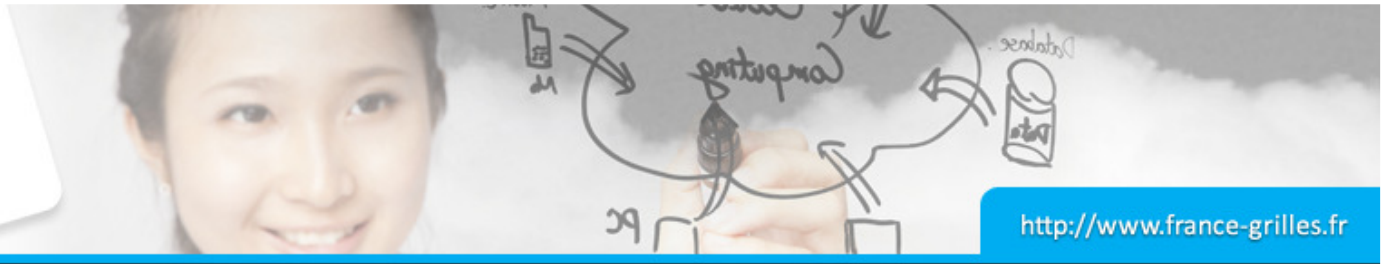

Next activities on clouds

- Collaboration with CSIC (Spain)
  - One person from CSIC to come to France for a year to work on cloud activities
- Federating French initiatives
  - Paris/GRIF and Stratuslab efforts
  - Cloud initiative and interests in Strasbourg/IPHC
  - Next cloud day in Lille - 11th May
- Contribute resources to EGI Cloud Federation TF

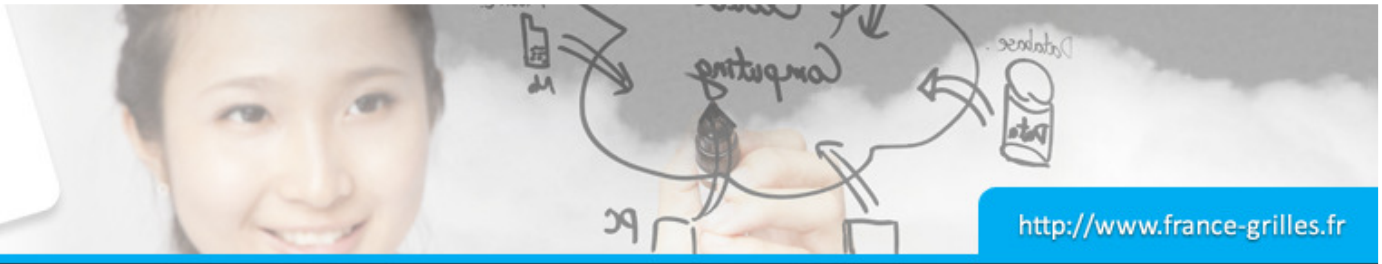

Conclusion on 2011 highlights and 2012 perspectives

- Addressed recommendations from last IAC
  - Realization of plans for a national VO
  - Initial description of our catalogue of services
  - Ongoing work on “easily understandable metrics to measure the performance and impact of IdG”
  - Laying foundation stones for a cloud infrastructure
- And now…
  - Time to think about longer term plans

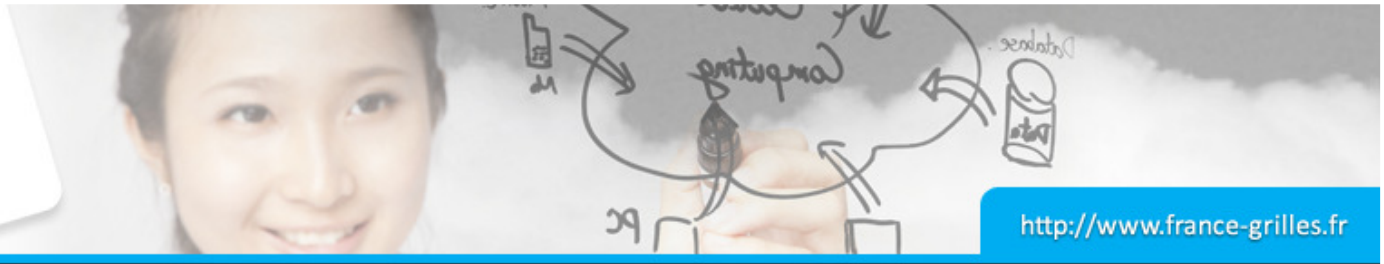

Outline

- Where are we ?
  - Some enlightening metrics
- Last year's highlights
  - Events and workshops
  - Catalogue of services
  - Internal tools and organization
  - Starting cloud initiatives
- Perspectives for the year to come
  - When Grid Operations meet Grid Users…
  - Getting more metrics: Methodology
  - Cloud perspectives
- Towards sustainability
  - Evaluating the cost and benefits of Grid Operations
  - Running national and global services beyond EGI-InSPIRE
  - What issues are we facing? How can we solve them?

YOU ARE HERE

The magic formula to successful operations

$$\frac{\frac{lim_{n\to\infty}\frac{2^{2n}(n!)^2 log7}{(2n)!\sqrt{n}}}{\int_0^\infty e^{-t^2}dt}}{[e \quad -e^{i\sum_{k=0}^{\infty}\frac{8}{(4k+1)(4k+3)}}]\int_0^\infty \frac{2t}{e^t-1}dt} = \mathbf{OK}$$

Grid
Non Grid

On all sites, Grid operations are tightly intricated with non grid operations

Cooperation and interdependencies

YESTERDAY

TODAY

TOMORROW ?

Cost of FG Operations (1) : Effort and manpower

- What is needed to play the "Grid Symphony"?
  - The orchestra: effort at site level
    - Estimation 30 FTE for site operations (incl. WLCG operations) over the current 18 production sites
  - Section leaders and conductor(s): effort at NGI level

NEEDED EFFORT					
Critical missions		Criticial and important missions		All missions	
h/week	FTE	h/week	FTE	h/week	FTE
186,5	**4,91**	**316,5**	**8,34**	**390**	**10,28**

    - 5 FTE for the minimum, 10FTE for the best *(see slide 12)*

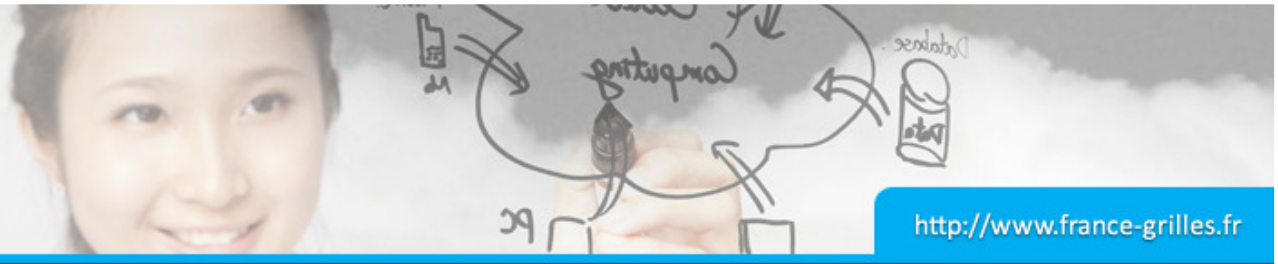

Cost of FG Operations (2) : Other costs

- Infrastructure costs
  - Hardware renewal and/or extension
  - Running grid resources and services (electricity, air conditioning, maintenance costs)
- Operating costs
  - Internal events (logistics and travel costs)
  - International involvement (travel costs , conference registration fees…)

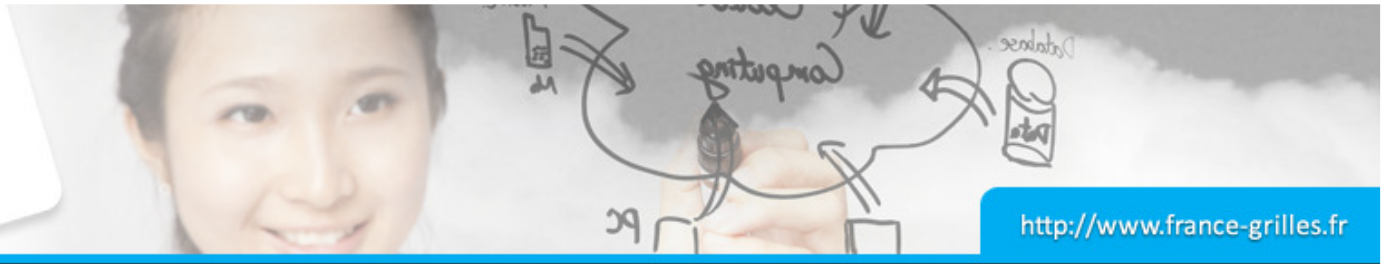

Benefits of FG Operations

- Direct benefits
  - Mutualization and sharing
    - Effort, investment and expertise : the "human Grid"
  - Coordination activities
    - m/w and service deployement, links to and between different projects
  - Central Grid services and functions
    - Monitoring, accounting, Certification Authority, Support
  - Grid security coordination
- Indirect benefits
  - Strategic collaborations (e.g. joint answers to project calls)
  - Visibility

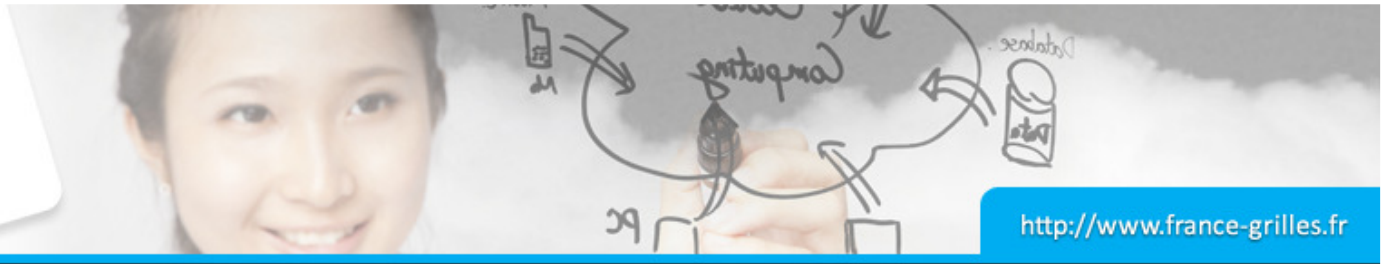

FG Operations : Current funding

- Who cover those costs today?
  - France Grilles - annual TGIR budget
    - 120k€ in 2012 for Operations
    - Hardware investments (e.g. 375k€ in 2012 on cloud infra.)
  - EGI-InSPIRE contribution for SA1 tasks
    - 1034 k€ for SA1 tasks over 4 years
  - Other projects funding source
    - WLCG, local grid projects, collaborations…
  - Sites/partners contributions to:
    - Effort
    - Travel and administrative costs for staff
    - Grid resources and services running costs

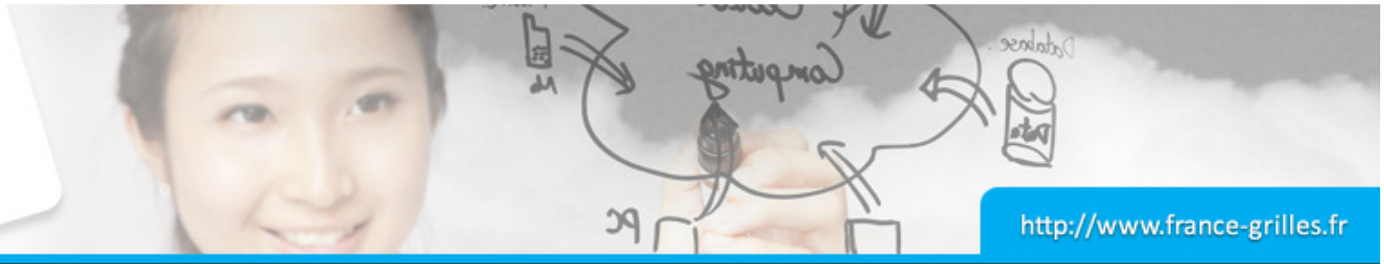

Beyond costs and benefits…

- How much of what we do would have to be done anyway, even if France Grilles didn't exist?
  - Probably most of site operations
    - Administration, operational support…
  - Most of high level expertise activities
    - Monitoring, accounting, Certification Authority, Security…

Sustainability questions

- Maintain living infrastructures
  - What happens if we don't have funding for :
    - The 35 needed FTEs?
    - Hardware & services maintenance costs?

- External dependencies
  - What if the tools we use stop being developed or maintained (e.g. EMI m/w, GOCDB, SAM-Nagios)?
  - What if we have to stop providing one of them to the Grid community (e.g. EGI Operations Portal)?

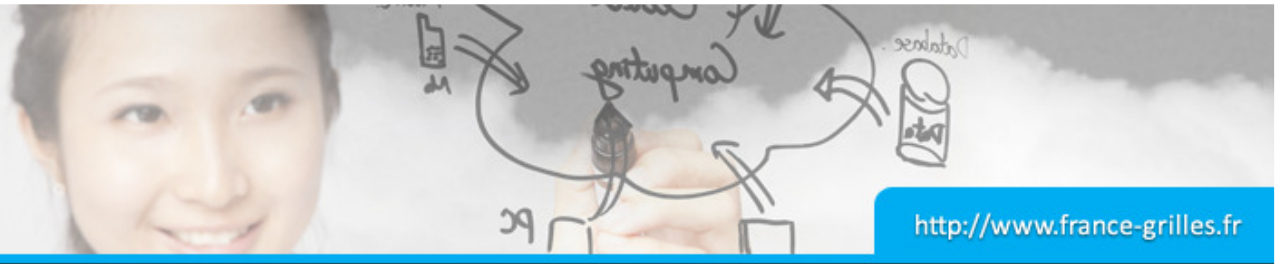

Some mitigation ideas (1)

- Consolidate our estimation of what we do
  - Go deeper in the costs/benefits analysis
  - Replace operations in the global picture

- Assess and clarify our “place in the landscape”
  - Formalize relations between France Grilles and
    - Its production sites
    - Its national partners
    - Its international partners

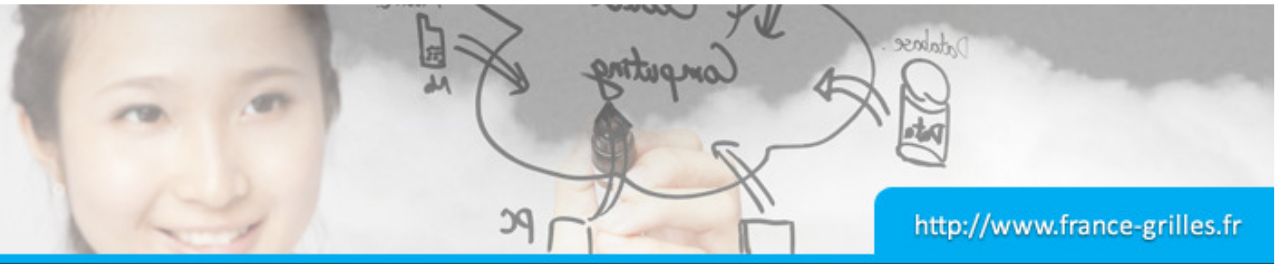

Some mitigation ideas (2)

- Build sustainability scenarios
  - Start from the hypotheses we can make on future funding, manpower, service availability etc.
  - Try and get ideas to turn each scenario into a form of success
- Get inspiration from elsewhere
  - What can Service Level Management teach us?
  - We are not alone… how are other NGIs doing?
  - Can we learn from the free software world?

Anyway…

- Whatever happens…

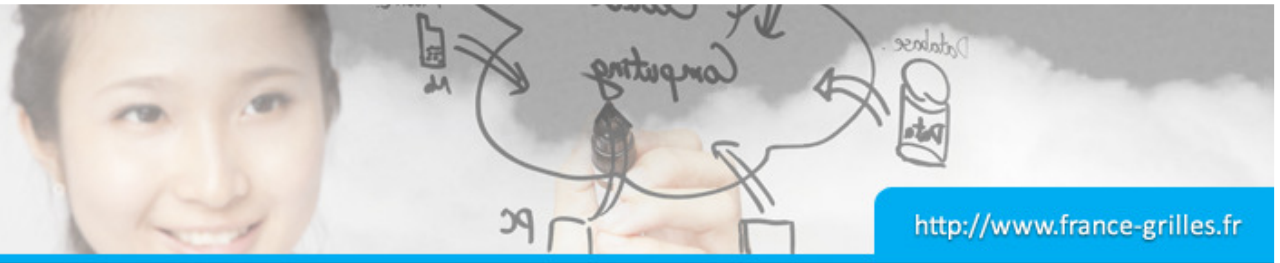

THANK YOU

Direction-technique@france-grilles.fr