Hepix/WLCG System Management WG

Alessandra Forti

Operations Workshop

14 June 2007

Layout

- Mandate
- WEB site
- Wiki
- Repositories
- Group
- Dissemination
- Conclusions

Mandate: Intro

- One of the problems observed (by EGEE and LCG) in providing a reliable grid service is the reliability of the local fabric services of participating sites.
- The SMWG should bring together the existing expertise in different area of fabric management to build a common repository of tools and knowledge for the benefit of HEP system managers' community.
- The idea is not to present all possible tools nor to create new ones, but to recommend specific tools for specific problems according to the best practices already in use at sites.
- Although this group is proposed in order to help improve grid sites reliability, the results should be useful to any site running similar local services.
- Two areas should be improved by the group: tools and documentation.

Mandate: Goals

- Improve overall level of grid site reliability, focusing on improving system management practices, sharing expertise, experience and tools
- Provide a repository
  - Management tools
  - Fabric monitoring sensors
  - HOWTOs
- Provide site manager input to requirements on grid monitoring and management tools
- Propose existing tools to the grid monitoring working group as solutions to general problems
- Produce a Grid Site Fabric Management cook-book
  - Recommend basic tools to cover essential practices, including security management
  - Discover what are common problems for sites and document how experienced sites solve them
  - Document collation of best practices for grid sites
- Point out holes in existing documentation sets
- Identify training needs
  - To be addressed in a workshop or by EGEE for example? We have been contacted and haven't replied yet.

Preliminary list of areas and tools

- System Management Areas
  - Filesystems: ext(2,3), XFS, NFS, AFS, dcache, DPM
  - Networking: Interfaces, IPs, Routers, Gateways, NAT
  - Databases: mysql, Oracle, ldap, gdbm
  - Processes: system, users monitoring
  - Servers: http, dhcp, dns, ldap, sendmail or other, sshd, (grid)ftp rfio
  - Batch systems: LSF, Torque, Maui, BQS, Sun Grid Engine, Condor
  - Security: login access pool accounts, certificates management and monitoring, non required services, ports list backups, monitoring(file systems, processes, networking), log files (grid services included)
  - .....
- Common Fabric Monitoring and Management Tools
  - Monitoring: Ganglia, Nagios, Ntop, Home grown, SAM, GridICE, Lemon
  - Management: Cfengine, Npaci rocks, Kickstart, Quattor
  - Security: iptables, rootkit, tripwire, nmap, ndiff, tcpdump, syslog, yummit
  - Grid Configuration: Yaim , Quattor

WEB site

- WEB site has been setup in Manchester
  - http://www.sysadmin.hep.ac.uk
- It's based on GridSite
  - http://www.gridsite.org
    - allows ACLs control based on x509 certificates
- The WEB site hosts
  - Wiki (Cookbook requested in the mandate)
  - Subversion repositories (sharing scripts)

Subversion Repositories

- http://www.sysadmin.hep.ac.uk/svnindex
- Integrated with GridSite
  - Read access is allowed to anyone
  - Write access based on certificates no need to create accounts but need to be added to the ACLs
  - Different repositories have different ACLs
    - Fabric-management (SMWG)
    - Fabric-monitoring (SMWG)
    - Grid-monitoring (GSWG)
      - Created, will be used soon.
    - Other. Creating a repo is very easy!

Subversion repositories (2)

- The tools should be management scripts or monitoring sensors written by sys admins to solve a local problem
  - However they should be generic enough to work at other sites
- Each script should have a banner containing the following information
  - Description
  - Author
  - Institute
  - Creation date
  - License
  - Repository version number
- Scripts not necessarily committed by the author
  - Always with their permission and license they want to use.
- There are currently 13 scripts in the repositories
  - We need more!

Wiki

- http://www.sysadmin.hep.ac.uk/wiki
- It is also integrated with GridSite
  - Accounts based on DN rather than user name and password.
- Simple rules to edit the wiki:
  - Each article should belong at least to one category to facilitate navigation and identification of the problem.
  - If the article contains a link to a script in the repositories it should belong to the “category scripts”
  - Each article or portion of article should bear the name and institute of the source if it is not the same as the page author.
    - For example if the text is extracted from a received email.

Wiki (2)

- Structure of categories is hierarchical with four top categories
  - Fabric management
  - Fabric monitoring
  - Best Practices (mostly basic and grid security)
  - Scripts to help navigate the repositories
- Subcategories are normally associated with a tool or one of the areas listed in a previous slide and then there are the articles.
  - Fabric Management (category) -> Cfengine (subcategory) -> Getting_started (article)
- Content at the moment:
  - there are 51 articles and 20 categories

Wiki(3)

- If good documentation is available some where else put just a pointer to the existing documentation.
  - Apply the minimum effort philosophy. For example Quattor page just points to the Quattor working group site after a small introduction.
  - But if someone wants to add an article with it's own experience can do it.
- Editing is currently done by me in a non systematic way.
  - Mostly assign articles to categories.
- However we used a wiki rather than writing a static document to avoid editing issues
  - Everyone should feel free to help writing an article or edit a stub.

SMWG Group

- Chairs:
  - Alessandra Forti (University of Manchester)
  - Michel Jouvin (LAL)
- Sent a call for participation to
  - HEPiX and all the T1s
- Mailing list:
  - wlcg-system-management-wg@cern.ch
  - 26 subscribers
- Meetings normally every fortnight the details are here:
  - http://www.sysadmin.hep.ac.uk/wiki/Meetings
  - Mainly to give updates about what people have done in the two weeks.
    - Haven't had one in a while will resume them after Stockholm.

SMWG Group (2)

- All the work is based on people volunteering to share
  - There are no dedicated people
- So there is no definition of group
  - Some people have only subscribed the mailing list
  - Some have subscribed the mailing list, participated to the meetings and done some work
  - Some people have acted as consultants and accepted their scripts to be distributed but are not on the mailing list nor come to the meetings
  - Some people have actually started editing the wiki with some stubs without even being in contact with any member of the group (i.e. mailing list subscribers)
- It's a start but it is not easy make people volunteer

Dissemination

- Sent an email to dcache user forum
  - They are the main users outside the group
- RSS feeds for OSCT point to some articles in the wiki
- Talks
  - HEPiX, UK HEP Sysman, GDB, Ops Workshop
- Should put a link from HEPiX WEB site
  - Michel is looking into it
  - Send an email to HEPiX mailing list when this is done
- Link from LCG WEB site?
  - Haven't discussed this with anyone yet
- Send an email to LCG-ROLLOUT
  - Was waiting to have a bit more content to convince people of the usefulness.

Conclusions

- There is a mandate
- There is a wiki
- There are repositories
- There is a group
- We need only people to contribute
- Questions?