OGF and Sausage

NGS Surgery, 24 Aug 2011

Jens

Open Grid Forum

OGF

Federated Identity Management

- “My Private Cloud” – fed access to Euc.
  - David Chadwick
  - Proposed API, for “cloud security” in PHP
    - authentication/identification, LoA
    - authorisation
    - “delegation by invitation”
    - Accounting “next”
- “Intercloud security” (GEYSERS projects)
  - Yuri Demchenko, UVA
- Contrail federation

Grids and Clouds: DCI (DCI?)

- Interest in Federated DCI – “standardise the use”
  - EGI is in, PRACE is not
  - XSEDE: sequel to TeraGrid (amalgamation of the two proposals)
  - GFFS: Global Federated FileSystem
  - Venus-C: “common roadmap”
  - Leasing resources/“cloudscaling”
  - Reference model contributions needed
  - Storage needed

Whither Then, Cloud Security?

- Previously: mostly show & tell
  - Useful, but mostly Usual Suspects(tm) s&t'ing
- How to progress usefully?
  - Would be useful to have milestones
  - A group which can produce documents
  - Needs commitment from people/projects
  - Draft charter
  - Scope: "DCI" fed security?
- What about delegation?
  - Role of EMI

EGI News

- “Need to use /etc”
  - Kebnekaise, Matterhorn, Monte Bianco
- Security components
  - VOMS, UVOS, ARGUS
  - Getting rid of GSI and moving to port type deleg.
- ESFRI role: “productisation”
  - CLARIN using iRODS
- WebDAV

EGI News

- Requirements
  - Must be important for the user communities
  - 1 month input, 11 months planning
- Storage Accounting Records (STAR)
- “More support for SAML & XACML”

Storage Accounting Records?

- GSM-CG : more detailed look (also https)
- OSG also producing SARs
  - Tracking data in flight
  - Tracking data at rest
- “Could we do “service accounting?””
- “Person from STFC who volunteered to contribute”
- Will UR replace non-service-discovery GLUE

CAOPS

- GFD.125: revision to make P-REC
  - Minor errata before public comments
- Private key protection guidelines
  - IGTF “encouraged” to write up guidelines
  - Opportunity for us to “manage” keys for users...
- RAT: no news
  - Needs communication test

FermiCloud

- Based on O'Neb.
  - Also evaluated Eucalyptus and Nimbus
- No stored secrets in image
- TODO:
  - Start in protected network
  - Wake dormant images for patches
- Pluggable X.509 contributed upstream
- Using XACML for AuZ

GLUE

- XML rendering online: 99% complete
- “GLUE2 not sufficient” for HPC
  - data staging protocols
  - parallel environments
  - session directory location
  - remote session directory access

Other Groups

- DRMAA
  - Working a lot “outside” OGF
- OCCI
  - Working a lot “outside” OGF
- Relations to OGC? (Contrail science case)

US Grids

- FutureGrid
  - Testbed modelled on Grid5000
  - Interop: not resource intensive
- Gordon:
  - Using flash extensively for storage
  - Apps similar to NGS
- Keeneland:
  - “Innovative HPC” based on GPU
  - Thrust. AccelerEyes. MAGMA (nearly all BLAS and LAPACK, multi-GPU in progress)
  - MAGMA: FERMI and TESLA have different arch; both supported. KEPLER to appear.
  - SHOC benchmark suites

Science Agency Use of Grids and Clouds:

SAUCG

~19 presentations

- How to move data between clouds if providers go out of business?
- OpenFlow: software-defined networking
  - LHCONE connectivity. DYNES data mover
- “Network as manageable resource”
- What makes a VO successful
  - derekweitzel.blogspot.com
- How PGI feed into clouds?
- How-to-measure-science (vs cost)?

SAUCG

- globusonline
  - Federated identity. RFT.
- Distributed data processing
  - Instrument data and metadata
- XSEDE
  - Genesis-II, Globus (XAUS (XD-Data)), Unicore
  - ByteIO, RNS, GFFS
  - (RNS="directory", EPR="inode")

SAUCG

- OPTIMIS project: SLAs, toolkits for optimising
  - Trust, risk, eco, cost, data security, legal
- SAGA+OCCI: case study
  - SAGA in front of: O'Neb (OCCI), EGI (GRAM), XD (BES), EC2 (ssh)
  - Mosaic cloud brokering
- SNIA plugfest: OCCI+CDMI
  - jclouds from SLA@SOI framework for talking to cloud
  - VenusC have CDMI client

SAUCG

- ORCA-BEN: Open Resource Control Arch.
  - “a kind of IaaS” (non-std), Hadoop on WAN
- FermiCloud
  - “Achieved significant savings in cooling and power”
  - Will move to SLAs, eg via WS-Agreement
  - Testbed: welcome new users
- Outsourcing studies
  - STAR nuclear physics (BNL): Use cloud as grids didn’t match work patterns
  - Bioinformatics: starting on desktop with VM, then moving to cloud deployment
  - “cloudinit.d”

SAUCG

- Contrail: ISIS, BADC, Genomics
- Orbiter project – work with RAL and DL
  - https://orbiter.txcorp.com/
- “Performance assured cloud storage”
  - PSpacer to control network bandwidth
  - PROBS: software based on OSD (reservation req’d)
- Focus on h’ware expenses (TU-Berlin)
  - No hardware in offices, all in machine room
  - bci.cit.tu-berlin.de

SAUCG

- “intercloud” – interoperable cloud
- Drivers behind push for cloud?
  - Non-human resource alloc?
  - Virtualisation?
  - Economy of scale based on underutilisation?
    - … but is this true in science?
  - “Inst’l resources better than nat’l”
  - Roadmaps

Summary of sorts

- Opportunities for NGS:
  - Collaborations at various levels
    - Grids and clouds
    - Users
  - Technology to play with
  - Stay with the (near-)production stuff