The trials and tribulations of the UF HPC Galaxy

Oleksandr Moskalenko, Ph.D.

UFIT Research Computing

Galaxy Caretakers Discussion on 09/20/12

Background

- UF HPC Center background and goals
- Galaxy goals

UF HPC Background

- Staff: 7 people
  - Systems and Operations – 3
  - Storage – 1
  - Applications – 1
  - Training – 1
  - Biocomputing support - 1
- Model:
  - Sustainability through share resource investments from faculty - $400 per 5 yr/PE, $0.02 billed rate
  - Support groups handling around $200E6 in grants
  - Non-investors get 8PEs for up to 24hrs per batch job

Galaxy goals

- Provide an easy to use interface for entry level analyses at first, expand the mission as the Galaxy's capabilities grow
- Provide a collaborative environment
- Provide a teaching environment
- Supplant the iNquiry instance to be retired from the Genetics Institute

UF Galaxy History and stats

- Deployment history
- Development and deployment setup
- Production configuration

History

- 1st prototype – KVM Virtual Machines – crash and burn
- 2nd prototype – real head node, Lustre FS, dedicated Torque/MOAB reservation
- 1st production setup – (10 months)
  - Head node (4cores, 16GB RAM)
  - Cluster (SGE – 4 12-core/96GB RAM nodes)
  - Lustre FS (60TB)
  - 1Gbit network, Infiniband from nodes to storage
  - Software provided by the modules system

Current production setup

- Number of active users (6mo) - ~150
- Size of the database/files – 5TB
- Jobs run: over 11000
- Hardware:
  - Head node – 4 cores, 16GB RAM (32GB upcoming)
  - Cluster
    - 7000 cores, 2GB/core average RAM
    - 120, 500GB RAM bigmem nodes (1TB upcoming)
    - Torque/MOAB PBS
  - Storage
    - Current – NexentaStor ZFS/NFS – slow
    - Upcoming – HA NexentaStor ZFS/NFS – fast

Deployment and development

- ▸ 3 head nodes: production, staging, development
  - ◦ Development:
    - • Up-to-date galaxy-central code
    - • Wrapper development
    - • Hacking on the galaxy core
    - • SQLite db
  - ◦ Staging:
    - • Pull galaxy-dist releases
    - • Pull and update a copy of the production database
    - • Clean and stabilize
    - • Set up tools and local modifications and wrappers
   - • Replace the production code

Production Configuration

‣ Authentication – Remote User
  ◦ Apache mod_auth_tkt, LDAP back-end.
‣ Hardware
  ◦ Head node – 4 cores, 16GB RAM
  ◦ Cluster
    • 700 nodes, 7000 cores, 2GB/core average RAM
    • 120, 500gb RAM bigmem nodes (upcoming 1TB)
    • Torque/MOAB PBS
  ◦ Storage
    • Current – NexentaStor ZFS/NFS – slow
    • Upcoming – HA NexentaStor ZFS/NFS – fast
  ◦ Software – modules and tool shed
  ◦ Database node – 4 core, 32GB RAM, PostgreSQL
  ◦ Local (campus) Tool Shed

Issues

- History
- Current
- Wish List

Initial Hurdles

▸ Torque - python_pbs memory leaks
▸ Head node overloaded
▸ Poor visualization capabilities
▸ No soft restart capability
▸ Upgrades – tool cleanup, maintaining hacks
▸ Providing consistent environment (drmaa.py / modules)
▸ No job accounting
▸ "I can haz moar wrapperzzz"

Pleaze sir…
I can has moar?

Early Hacks

- Watcher cron job for python_pbs
- Run **all** jobs on the cluster, even the upload
- Soft restart init scripts
- Use two RCS – mercurial and git to manage upgrades
- Hacked drmaa.py to load modules
- Unified the tools used on the cluster and in Galaxy via the environment modules system
- Unified reference databases
- Set up big data import

Finding Help

▸ Galaxy wiki is a great resource for experienced minds
  ▸ http://gmod.827538.n3.nabble.com/Galaxy-Development-f815885.html (galaxy-dev list archive on Gmod) and

  ▸ http://gmod.827538.n3.nabble.com/Galaxy-Users-f815892i36.html (galaxy-user list archive on Gmod) are invaluable

▸ http://seqanswers.com/ - treasure trove of information on tools when a bug report calls.

▸ Read the code

▸ IRC (#galaxyproject) – mostly quiet, but getting a core developer to commit a new feature to galaxy-central after a 3 minute discussion is priceless when it happens

▸ Galaxy Community Conference

▸ Blogs

User Issue Report Categories

▸ Inappropriate job resource requests hard-coded into the tool runner URIs

▸ Tool failures because of bad data or tool bugs

▸ Poor understanding of what tools are needed

▸ Help with analyses

▸ Reference dataset issues (dbkey)

Last Upgrade

- Sep 2012 upgrade
- Wish List

The "Big Upgrade" – Sep 2012

- Jobs must run under the real user's id
  - Very powerful feature, but opens many cans of worms
- Move from SGE to Torque/MOAB
  - No green banner of death
  - Strict resource request handling
  - Finding a drmaa library that works - http://apps.man.poznan.pl/trac/pbs-drmaa
- Introduction of the Tool Shed
  - Cultivate new wrapper contributors
  - Very unstable at the moment
  - Crucial divergence from the modules in version handling
- Dynamic job runner
  - Almost unlimited possibilities in with the real userid jobs

The Wish List

▸ Universal job resource request interface mechanism

▸ More flexible output file handling

▸ No more hard-coding tool arguments

▸ Easier access to core developers

▸ Reference dataset handling automation

▸ More documentation and more capabilities in the tool definition file logic

▸ More published workflows and docs

Questions?