ASL ‘12
Project infrastructure

Thursday, September 20, 2012

Who we are:

- Jana Giceva
- Markus Pilman
- Lucas Braun
- Georgios Giannikis
- Tudor Salomie
- Simon Loesing

Group Registration

- Groups of 3 students
- Sheets are on the front desk, come up and write your name (at the **end** of the session)
- People that have not registered on the sheet but have registered for the course will be randomly assigned to groups

The group distribution and TA assignment will be posted online once groups are finalized. (i.e., Monday?).

About the project

- Implementing a distributed messaging system
- 3 milestones: coding, benchmarking, modeling
  - Every milestone is graded
  - All milestones have to be submitted
  - Project is part of final grade
- Java and PostgreSQL
  - Mandatory, no other PL or DBMS
- For deliverables
  - NO: external libs (except JDBC driver), RMI, JPA
- For development
  - YES: scripting, sed/awk/grep, cron jobs, ssh, python, etc.

Infrastructure

- For testing and benchmarking you need multiple servers w/ dedicated access
  - Running on one laptop is **NOT OK**
  - Development on laptop is **OK**
- For benchmarking there are 2 options
  - Dryad cluster in ETHz
  - Amazon EC2 in the cloud

Dryad cluster

- 16 machines
  - dryad01 to dryad16
  - 8 cores, 16GB or RAM and 2 x 1TB harddrives
  - only accessible from ETH network (or w/ VPN)

Dryad cluster - access

- Each group is assigned an account (non-ethz)
  - non-root accounts, local on each dryad machine
  - access is **only** over ssh/scp
- Account name like userXX
- Access & Authorization is enforced
- Server time split in 4h slots that will be published on course web-page

Every student needs to send us his/her ssh pub-key (no access otherwise); use ssh-keygen to generate it

Dryad cluster – NFS

- Store experimental output on /mnt/asl/
  - Create folder with group name
  - Keep it below 1GB per group
  - Keep it below 1000 files per group
- Do not rely on data on the local HDD to persist
- The /mnt/asl is accessible from sgv-asl-01.ethz.ch when you do not have a booked slot
- sgv-asl-01 should only be used to scp data in/out
  - Virtualized environment (probably slower than your laptop)

Amazon EC2

- Virtual machines in a Amazon Datacenter
  - Different Instance types (slow – super fast)
  - Existing free images for all linux distributions
  - Managed from a Web-Console
  - Access to VMs over SSH w/ root access
- You have to create an AWS Account
  - Credit Card Needed!
  - Each new account gets a free usage tier (750h on micro instances / month)
  - We may get some educational vouchers (100$/group)
- Well documented: check www.aws.amazon.com

SVN repos

- All code that you write should be kept in the SVN repo that we provide:

`https://svn.inf.ethz.ch/svn/systems/asl12/trunk/userXX`

  - userXX is the one assigned to your group
- Login with N-ETHz account
- Submissions will be done via SVN
- On deadline we check-out for each group:
  - `…/userXX/milestone01/`
  - `…/userXX/milestone02/`
  - `…/userXX/milestone03/`
- **Only content in these folders will be considered for the submission**

Next steps (I)

- Finalize groups
- Send us your ssh pub-keys:
  - mailto: sg-asl [at] lists [dot] inf [dot] ethz [dot] ch
  - subject: “ssh pub-key”
  - from: <nethz-account> [at] student.ethz.ch
- Test SVN access & login to dryads
- Build PostgreSQL from source

Next steps (II)

- Meeting w/ your TA in assigned classroom (will be online on the course web-page by Monday)
- 27th Sept:
  - each group does a 10 min presentation
  - database schema for messaging system
  - interfaces:
    - middleware –to– database
    - client –to– middleware