Introduction to Grid’5000

Chuyuan Li
22/09/2021

What is Grid’5000 and why would you use it?

This is a large-scale and flexible testbed for experiment-driven research.

We mainly interested in its large amount of resources:

- when you want to run a GPU-required machine learning task but you don't have GPU in your own computer
- when you run a time-consuming calculation and wish not to occupy 90% of your CPU all the time, *etc*.

For detailed description, refer to this link

Outline

- Part 1: 22/09 16H-18H
  - Get an account of Grid’5000
  - Connection with SSH key
  - Basic concepts (cluster, node, host, core...)
  - File/folder transfer
- Part 2: 01/10 16H-18H
  - Resources visualisation
  - Resources reservation and management with OAR
  - *TBD*

Before we start...

- Do you have an account and activated?
  - Open a terminal
  - type: ssh login@access.grid5000.fr
- Sites you will frequent:
  - Get started: https://www.grid5000.fr/w/Getting_Started
  - UMS: https://api.grid5000.fr/stable/users/
  -

Basic SSH

Authenticating

- SSH = Secure SHell
- Standard network protocol and service, establish a secure communication channel between 2 machines
- Relies on cryptography
- Public-key authentication
  - general idea: asymmetric cryptography
  - **public key** is used to **encrypt** something
  - only the **private key** can **decrypt** it
  - user owns a private key, stored on the local machine
  - server has the public key corresponding to the private key
  - authentication = server proves that you own that private key

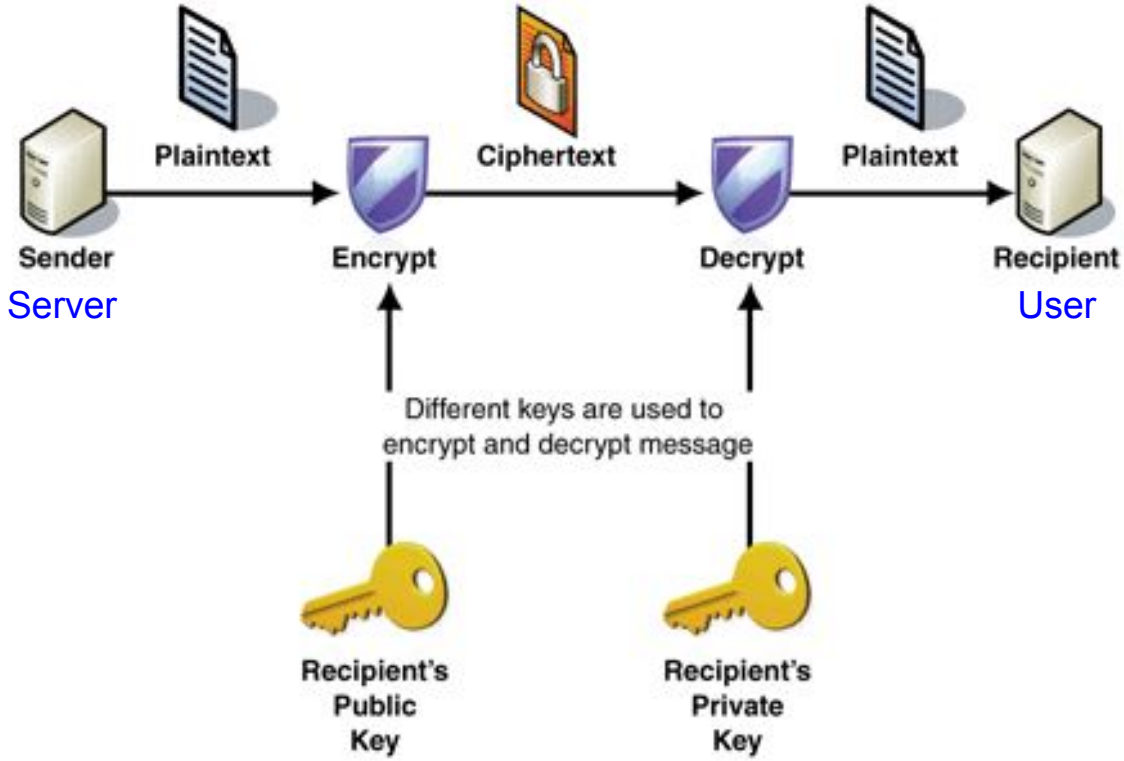
Plaintext
Ciphertext
Plaintext
Sender
Encrypt
Decrypt
Recipient
Server
User
Different keys are used to
encrypt and decrypt message
Recipient's
Public
Key
Recipient's
Private
Key

SSH basic usage

- Most basic usage: get shell access on a remote machine
- Many advanced usages:
  - data transfer (`scp`, `sftp`, `rsync`)
  - connect to specific services (such as Git or SVN servers), *etc*.
- Connecting to a remote server
  - $ ssh login@remote-server
  - this provides a shell on remote-server
- Copying data
  - $ scp local-file login@remote-server:remote-directory/
  - $ scp login@remote-server:remote-dir/file local-dir/
  - $ rsync -avzP localdir login@server:path-to-rem-dir/
- Know more here and here

Big picture

Concept map

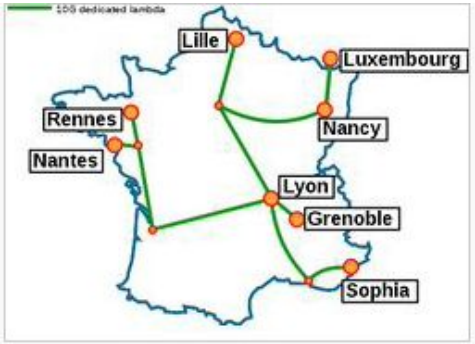
10G dedicated lambda
Lille
Luxembourg
Rennes
Nancy
Nantes
Lyon
Grenoble
Sophia

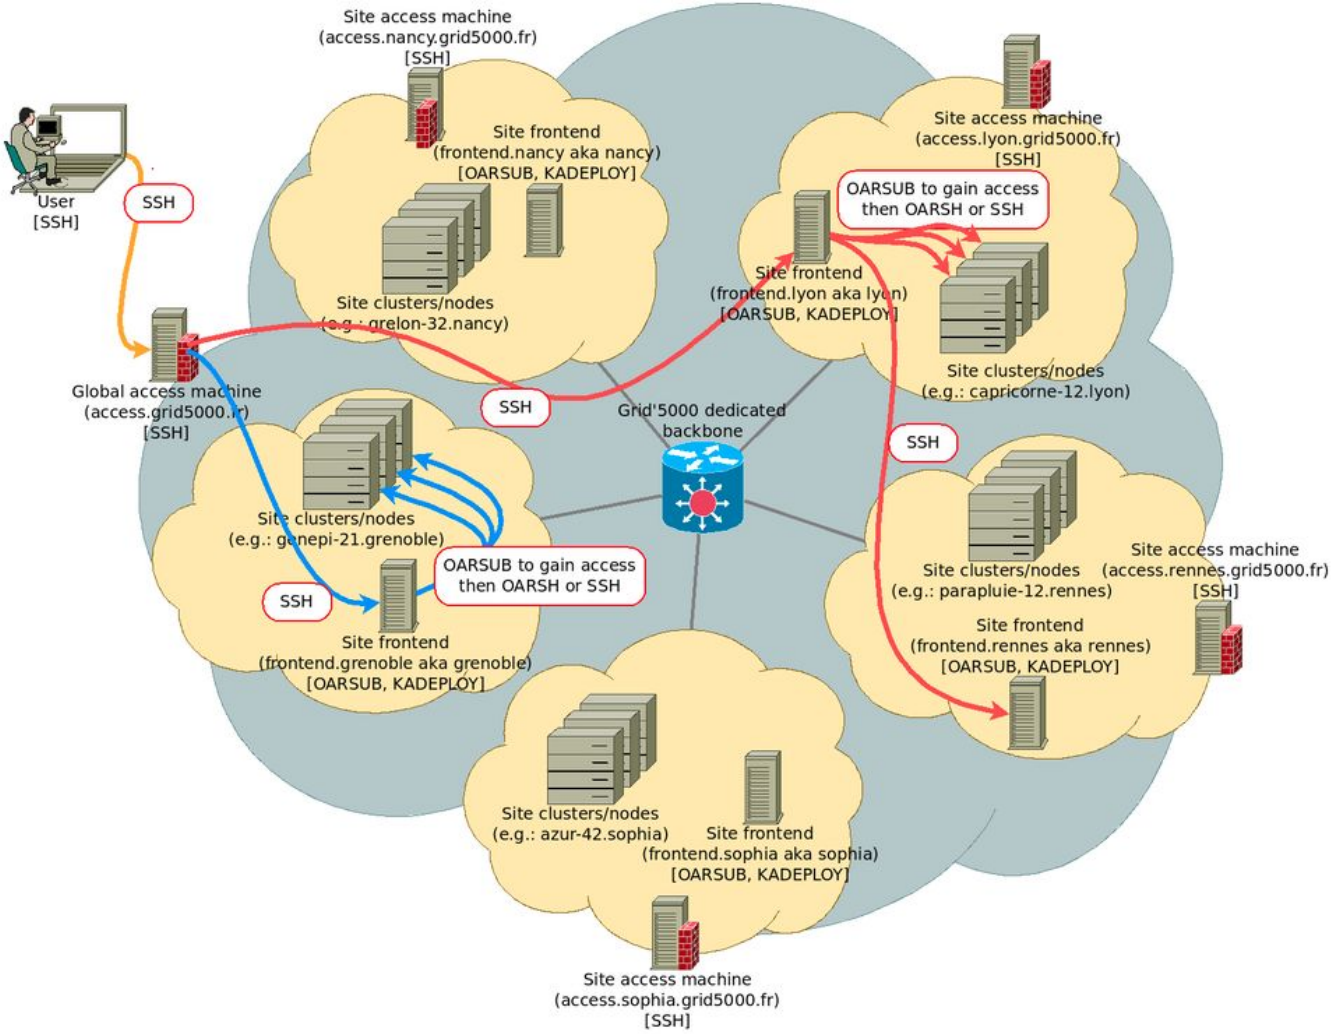
User
[SSH]
SSH
Global access machine
(access.grid5000.fr)
[SSH]
Site access machine
(access.nancy.grid5000.fr)
[SSH]
Site frontend
(frontend.nancy aka nancy)
[OARSUB, KADEPLOY]
Site clusters/nodes
(e.g.: grelon-32.nancy)
SSH
Grid'5000 dedicated
backbone
Site access machine
(access.lyon.grid5000.fr)
[SSH]
OARSUB to gain access
then OARSH or SSH
Site frontend
(frontend.lyon aka lyon)
[OARSUB, KADEPLOY]
Site clusters/nodes
(e.g.: capricorne-12.lyon)
SSH
Site clusters/nodes
(e.g.: genepi-21.grenoble)
OARSUB to gain access
then OARSH or SSH
SSH
Site frontend
(frontend.grenoble aka grenoble)
[OARSUB, KADEPLOY]
Site clusters/nodes
(e.g.: parapluie-12.rennes)
Site access machine
(access.rennes.grid5000.fr)
[SSH]
Site frontend
(frontend.rennes aka rennes)
[OARSUB, KADEPLOY]
Site clusters/nodes
(e.g.: azur-42.sophia)
Site frontend
(frontend.sophia aka sophia)
[OARSUB, KADEPLOY]
Site access machine
(access.sophia.grid5000.fr)
[SSH]

Concept map

- cluster
- node
- host

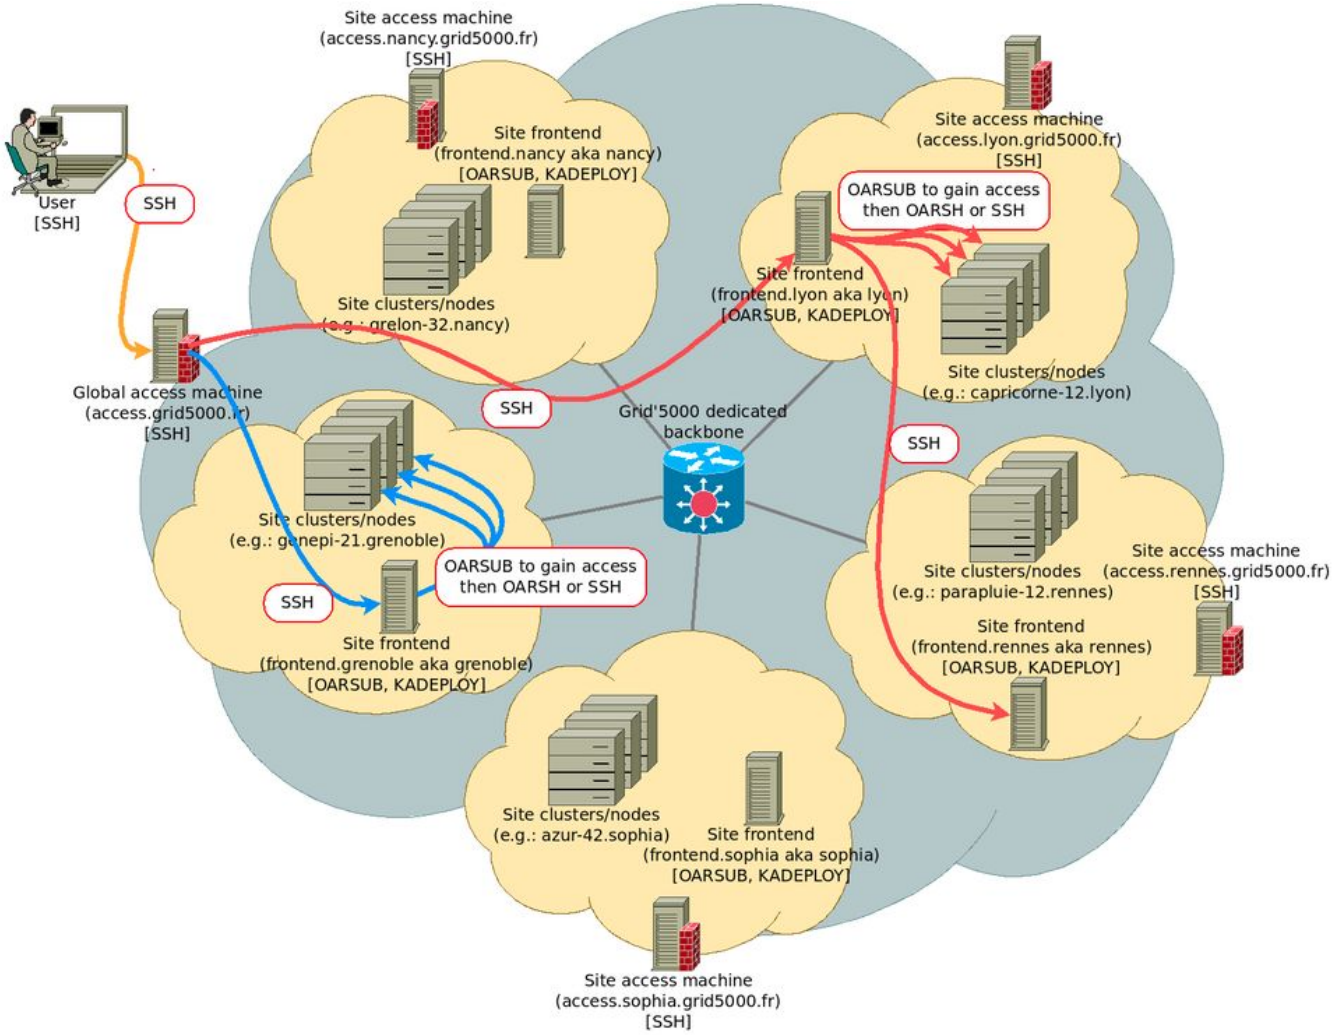

Hardware in Nancy site

13 clusters, 374 nodes, 7784 cores, 323.3 TFLOPS

Cluster	Access Condition	Date of arrival	Nodes	CPU	Cores	Memory	Storage	Network	Accelerators
graffiti	**production** queue	2019-06-07	13	2 x Intel Xeon Silver 4110	8 cores/CPU	128 GiB	**479 GB SSD**	10 Gbps	4 x Nvidia RTX 2080 Ti
graoully	**production** queue	2016-01-04	16	2 x Intel Xeon E5-2630 v3	8 cores/CPU	128 GiB	**1 x 600 GB HDD** + 1 x 600 GB HDD	10 Gbps + 56 Gbps InfiniBand	
graphique	**production** queue	2015-05-12	6	2 x Intel Xeon E5-2620 v3	6 cores/CPU	64 GiB	**299 GB HDD**	10 Gbps + 56 Gbps InfiniBand	1: 2 x Nvidia Titan Black [2-6]: 2 x Nvidia GTX 980
graphite		2013-12-05	4	2 x Intel Xeon E5-2650	8 cores/CPU	256 GiB	**1 x 300 GB SSD** + 1 x 300 GB SSD	10 Gbps + 56 Gbps InfiniBand	Intel Xeon Phi 7120P
grappe	**production** queue	2020-08-20	16	2 x Intel Xeon Gold 5218R	20 cores/CPU	96 GiB	**480 GB SSD** + 8.0 TB HDD*	25 Gbps	
grcinq	**production** queue	2013-04-09	47	2 x Intel Xeon E5-2650	8 cores/CPU	64 GiB	**1.0 TB HDD**	1 Gbps + 56 Gbps InfiniBand	
grele	**production** queue	2017-06-26	14	2 x Intel Xeon E5-2650 v4	12 cores/CPU	128 GiB	**1 x 299 GB HDD** + 1 x 299 GB HDD	10 Gbps + 100 Gbps Omni-Path	2 x Nvidia GTX 1080 Ti
grimani	**production** queue	2016-08-30	6	2 x Intel Xeon E5-2603 v3	6 cores/CPU	64 GiB	**1.0 TB HDD**	10 Gbps + 100 Gbps Omni-Path	2 x Nvidia Tesla K40M
grimoire		2016-01-22	8	2 x Intel Xeon E5-2630 v3	8 cores/CPU	128 GiB	**600 GB HDD** + 4 x 600 GB HDD* + 200 GB SSD*	4 x 10 Gbps + 56 Gbps InfiniBand	
grisou		2016-01-04	51	2 x Intel Xeon E5-2630 v3	8 cores/CPU	128 GiB	**1 x 600 GB HDD** + 1 x 600 GB HDD	[1-48]: 1 Gbps + 4 x 10 Gbps 49: 4 x 10 Gbps [50-51]: 4 x 10 Gbps + 56 Gbps InfiniBand	
gros		2019-09-04	124	Intel Xeon Gold 5220	18 cores/CPU	96 GiB	**480 GB SSD** + 960 GB SSD*	2 x 25 Gbps	
grue	**production** queue	2019-11-25	5	2 x AMD EPYC 7351	16 cores/CPU	128 GiB	**479 GB SSD**	10 Gbps	4 x Nvidia Tesla T4
grvingt	**production** queue	2018-04-11	64	2 x Intel Xeon Gold 6130	16 cores/CPU	192 GiB	**1.0 TB HDD**	10 Gbps + 100 Gbps Omni-Path	

Link: https://www.grid5000.fr/w/Nancy:Hardware

Queues and Usage Policy

- *Default* queue
  - Daytime is dedicated to smaller-scale experiments
  - Large-scale jobs must be executed during nights or weekends
  - generally, using advance reservations
  - **Read carefully the rules** in case of violation of usage
- Production queue
  - Smaller set of resources
  - Only in Nancy site
  - More suited to long-running, non-interactive jobs
- More information, ref to UsagePolicy

Exercise (1) How many hours can I reserve?

According to Usage Policy for Default queue:

> *Between 09:00 and 19:00 during working days (Monday to Friday, excluding public holidays in France), you should not use more than the equivalent of 2 hours on all the cores of the cluster during a given day (e.g. on a 64 bi-processor (quad core) cluster, you should not use more than (2 hours)*(2 CPU)*(4 cores)*(64 nodes)= 1024 core.hours).*

For cluster grimoire, if i want to reserve 32 cores for a task, how long can I reserve?

Cluster	Access Condition	Date of arrival	Nodes	CPU	Cores	Memory	Storage	Network	Accelerators
grimoire		2016-01-22	8	2 x Intel Xeon E5-2630 v3	8 cores/CPU	128 GiB	**600 GB HDD** + 4 x 600 GB HDD* + 200 GB SSD*	4 x 10 Gbps (SR-IOV) + 56 Gbps InfiniBand	

Queues and Usage Policy

- discover daily allowance with:

  `usagepolicycheck -l [--sites site1,sites2]`
- check the jobs that have been counted using:

  usagepolicycheck -v --start '2021-06-01 11:00:24 +0200' --end '2021-06-20 10:00:24 +0100'

First connection

Connecting and moving around

- Basic steps to get in a site:
  - open a terminal
  - connect to access machine: `outside: ssh login@access.grid5000.fr`
  - specify a site: `access: ssh site`
  - put in your password
  - then we can view machine list in this site

Connecting and moving around

- Basic steps to get in a site:
  - connect to access machine: `outside: ssh login@access.grid5000.fr`

```
[chuyli@lisa27:~$ ssh cli@access.grid5000.fr
Linux access-north 4.9.0-12-amd64 #1 SMP Debian 4.9.210-1+deb9u1 (2020-06-07) x86_64
----- Grid'5000 - access-north.grid5000.fr -----

Welcome to Grid'5000

** Connect to a site:
$ ssh {grenoble,luxembourg,lyon,nancy,nantes,rennes,sophia}

** Useful links:
 - account management (password change): https://api.grid5000.fr/ui/account
 - homepage: https://www.grid5000.fr/mediawiki/index.php/Category:Portal:User
 - charter : https://www.grid5000.fr/mediawiki/index.php/Grid5000:UserCharter
 - support : https://www.grid5000.fr/mediawiki/index.php/Support

** Data on access.grid5000.fr :
 - your home directory on access machines (access-north and access-south)
   is not synchronized and should not be use to store data.
 - please use ssh forwarding to send data directly to sites or
 - (outside) $ scp files login@access.grid5000.fr:reims/ using the nfs mount point in your home
```

  - specify a site: `access: ssh nancy`
  - Move around with ssh: global access -> nancy -> grenoble -> rennes…
  - `exit` to get out

Tip: use SSH ProxyCommand

- In ~/.ssh/config:

```
Host g5k
  User USERNAME
  Hostname access.grid5000.fr
  ForwardAgent no
Host *.g5k
  User USERNAME
  ProxyCommand ssh g5k -W "$(basename %h .g5k):%p"
  ForwardAgent no
```

- Only works is login shell is bash, otherwise need to adapt it
- For windows users, different options see: https://www.grid5000.fr/w/SSH#Windows_users
- Connect to any Grid5k node in one command
    - $ ssh nancy.g5k
    - $ ssh lyon.g5k

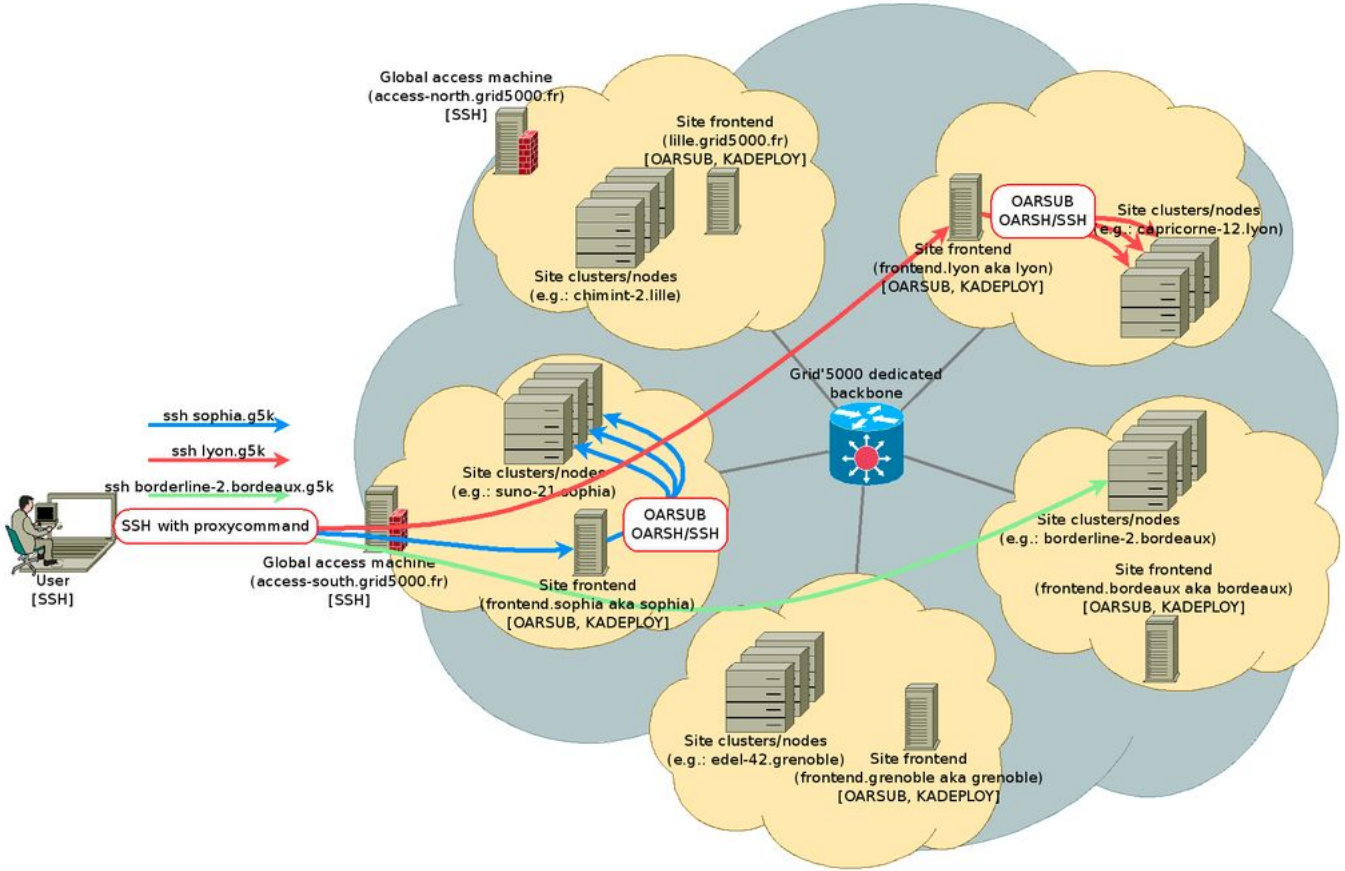
Global access machine
(access-north.grid5000.fr)
[SSH]
Site frontend
(lille.grid5000.fr)
[OARSUB, KADEPLOY]
Site clusters/nodes
(e.g.: chimint-2.lille)
OARSUB
OARSH/SSH
Site clusters/nodes
(e.g.: capricorne-12.lyon)
Site frontend
(frontend.lyon aka lyon)
[OARSUB, KADEPLOY]
Grid'5000 dedicated
backbone
ssh sophia.g5k
ssh lyon.g5k
ssh borderline-2.bordeaux.g5k
SSH with proxycommand
User
[SSH]
Global access machine
(access-south.grid5000.fr)
[SSH]
Site clusters/nodes
(e.g.: suno-21.sophia)
OARSUB
OARSH/SSH
Site frontend
(frontend.sophia aka sophia)
[OARSUB, KADEPLOY]
Site clusters/nodes
(e.g.: borderline-2.bordeaux)
Site frontend
(frontend.bordeaux aka bordeaux)
[OARSUB, KADEPLOY]
Site clusters/nodes
(e.g.: edel-42.grenoble)
Site frontend
(frontend.grenoble aka grenoble)
[OARSUB, KADEPLOY]

Transferring files to/from Grid’5000

- no BACKUP in g5k, so make sure your important files are stored somewhere outside
- In each site, by default 25 GiB storage
  - If needed, can demand for more space
  - *manage account -> homedir quotas -> request quota extension*
- ProxyCommand works with everything SSH-based
  - scp, sftp, rsync
- Prefer **rsync** than scp
  - Pipelined file transfers
  - More efficient on networks with large BDP (bandwidth * latency)

Transferring files to/from Grid'5000

- scp
  - Copy file from local to remote:
    - `scp local_file remote_username@remote_ip:remote_file`
  - Copy folder from local to remote:
    - `scp -r local_folder remote_username@remote_ip:remote_folder`
  - Copy file from remote to local:
    - `scp remote_username@remote_ip:remote_file local_file`
  - Copy folder from remote to local:
    - `scp -r remote_username@remote_ip:remote_folder local_folder`
- Example
  - `local: $ scp -r /Users/chuyli/g5k_tuto/ cli@nancy.g5k:/home/cli/g5ktuto`
  - `local: $ scp -r /Users/chuyli/g5k_tuto/ cli@access.grid5000.fr:nancy/g5ktuto`
  - `local: $ scp cli@nancy.g5k:/home/cli/g5ktuto/show1.sh /Users/chuyli/g5k_tuto/`

Transferring files to/from Grid’5000

- rsync
  - Copy folder from local to remote:
    - `rsync -avzP local_folder remote_username@remote_ip:remote_folder`
  - Example:
    - `local: $ rsync -avzP /Users/chuyli/g5k_tuto cli@nancy.g5k:/home/cli/`
    - `local: $ rsync -avzP /Users/chuyli/g5k_tuto/ cli@nancy.g5k:/home/cli/`
    - Mind the difference between *local_folder* and *local_folder/*

  - Much faster than scp for large files, recommend for folder transfer
  - Syntaxe more complicated
  - To know more, check official link rsync

Transfer of 120 files (total : 2.1 MB) with SCP and Rsync
Bandwidth and Latency controlled using network emulator

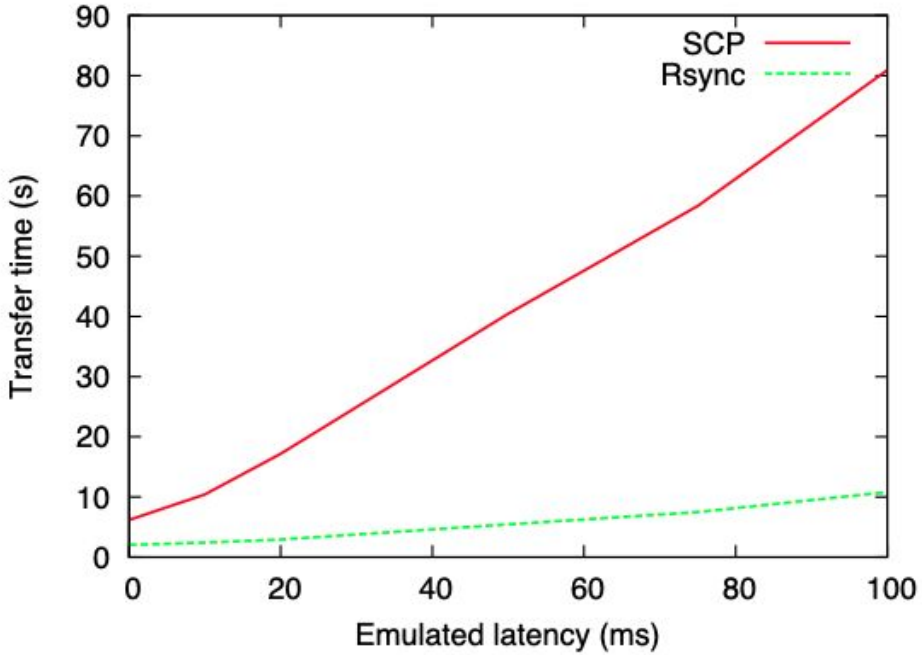

Exercise (2)

- Transfer from local to remote a folder called g5ktest/ with 2 files inside with rsync, remote folder should contain the same folder
- Transfer from remote to local a file called *remote2local.txt* with scp