



# SNIC and UPPMAX

Compute, Storage, and Cloud

Marcus Lundberg

[marcus.lundberg@uppmax.uu.se](mailto:marcus.lundberg@uppmax.uu.se)



**SNIC**

# Overview of talk

- Organisational orienteering
- Available resources (systems)
- Overview of the environment
- **Break**
- Practical stuff

# Swedish HPC Resources

- **SNIC** — Swedish National Infrastructure for Computing
  - Free to use — funded by VR + university consortium
  - From 2023 — replaced by NAIS
- **UPPMAX** — HPC center with SNIC resources at UU
- **SNIC-SENS** — Secure system for sensitive personal data
- **SNIC Science Cloud (SSC)** — SNIC's cloud
- **SNIC Application Experts** — advanced user support




# SNIC Clusters

- **Rackham (UPPMAX)** - Large storage w/ backup
- **Bianca (UPPMAX)** - Sensitive data
- **Dardel (PDC)** - Large system, highly parallel codes, general-purpose codes, to be expanded with GPU nodes
- **Tetralith (NSC)** - Large general-purpose system, has GPU nodes
- **Alvis (C3SE), BerzeLiUs (NSC)** - Systems for AI/ML programs

# Snowy

- UU-funded cluster
- 200 nodes in total (16 cores, 128 GB RAM)
- 50 nodes with 1 Nvidia T4 GPU
- Log in via Rackham
- Extra long jobs



# Alvis (AIPI)

- Resource dedicated to Artificial Intelligence and Machine Learning research
  - Quite large, with varied hardware
  - Popular datasets and a lot of software available
  - Pip, Singularity, and Conda available to provision your own software.
- <https://portal.c3se.chalmers.se/public/root/>

# BerzeLiUs

- Resource dedicated to Artificial Intelligence and Machine Learning research
  - NVIDIA SuperPOD with 60 DGX-A100 nodes
- <https://www.nsc.liu.se/systems/berzelius/>

# SNIC Science Cloud

- Offers Infrastructure-as-a-Service
  - Open to provisioning PaaS and SaaS as well, needs pilot use cases
- Divided into “regions”, each with the same OpenStack-based environment
- Most useful for relatively small-scale projects, e.g. development of a cloud-based application
- UPPMAX Cloud is now equipped with 24 NVIDIA GPUs (20xA2 and 4xT4)!



# What SNIC can/can't be used for

- **Do:**
  - Analyse data
  - Store data while you're working with it
- **Do not:**
  - Use to (only) store/backup/move a dataset
  - Run a business



# Applying for SNIC project

- Projects are managed in SUPR ([supr.snic.se](https://supr.snic.se))
- Resources are granted to projects
- Project membership determines which files are readable and how much compute power is available
- <https://www.uppmax.uu.se/support/getting-started/applying-for-projects/>

# Scale of project

	Compute (core-hours/ month)	Storage (TB)	SENS
Small	< 10,000	< 10	< 10 kch/m < 20 TB
Medium	< 200,000	< 100	< 100 kch/m < 150 TB
Large	> 200,000	> 100	> 100 kch/m > 150 TB

# The cost?

- Free to you but...
- Rackham's extension cost (very roughly, only initial investment):
  - 375 kr/TB per year
  - 0.1 kr/core-hour
- A typical PhD-student's project (1 TB, 1000 core-hours/month for four years): 6,300 kr.
- Large molecular dynamics simulations, using 200,000 core-hours/month: 240,000 kr/year.

# Using SNIC/UPPMAX

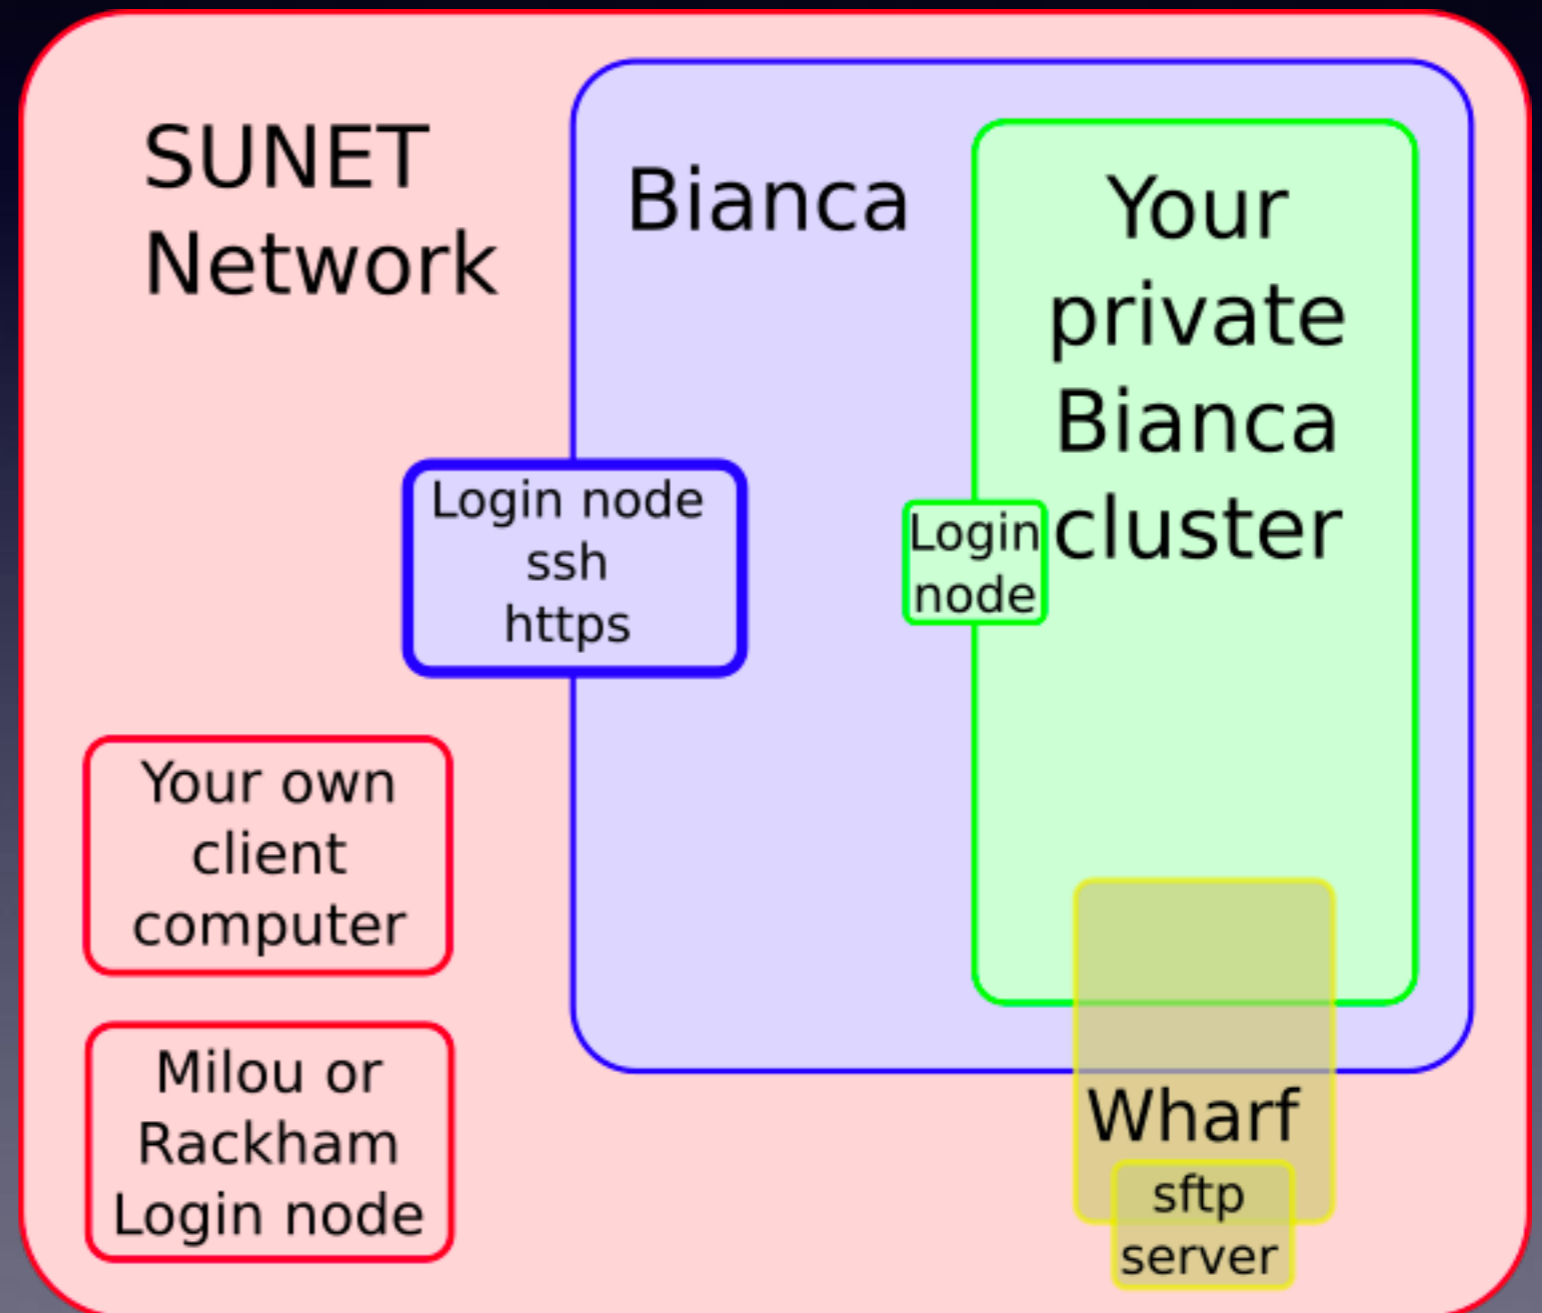
- Linux command-line
- Graphical remote Linux desktop
- Training available via SNIC, UPPMAX
  - <https://www.uppmax.uu.se/support/courses-and-workshops/>

# Sensitive personal data

- What is (or might be) sensitive personal data?
- Why do you need additional technical measures?

# Bianca (SNIC SENS)

- “Lätt att göra rätt”
- A secure HPC environment
- Collaboration-friendly



# Software at UPPMAX

- A huge software library is installed and maintained
- Additional software can be requested, but be aware this can take a while
- **Most software can be installed by yourself**

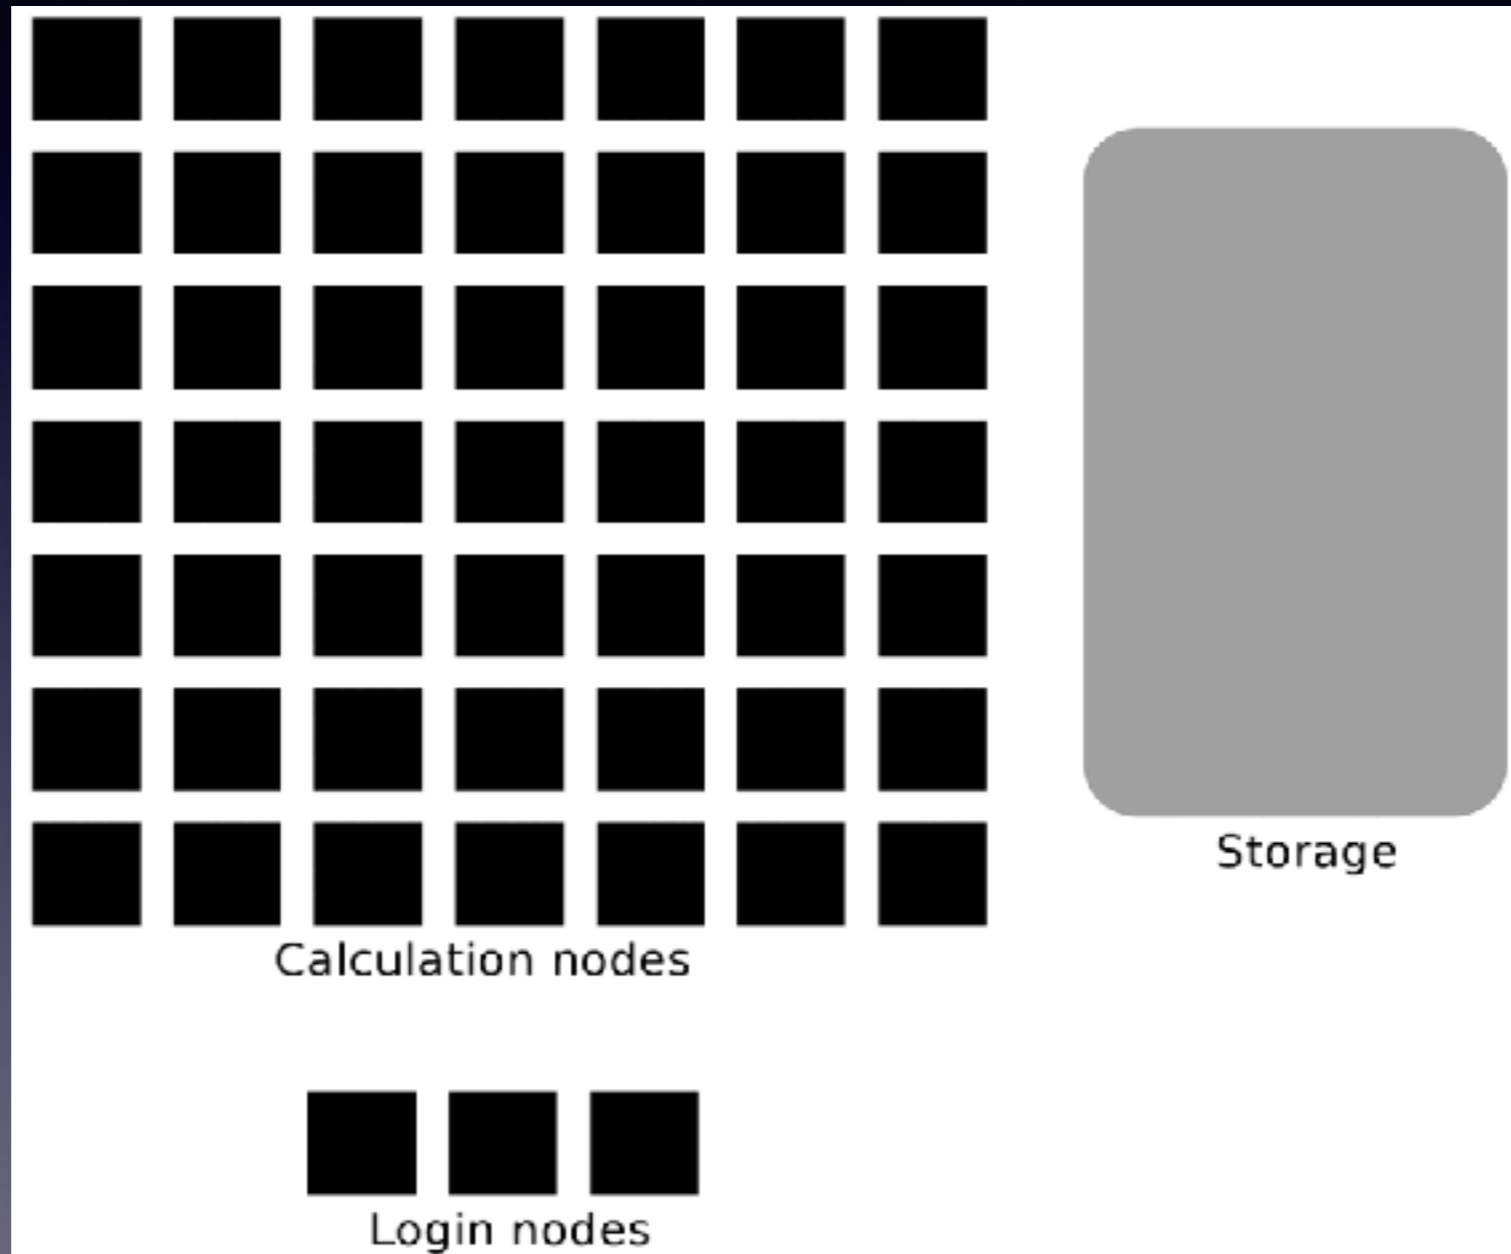


# Finding software

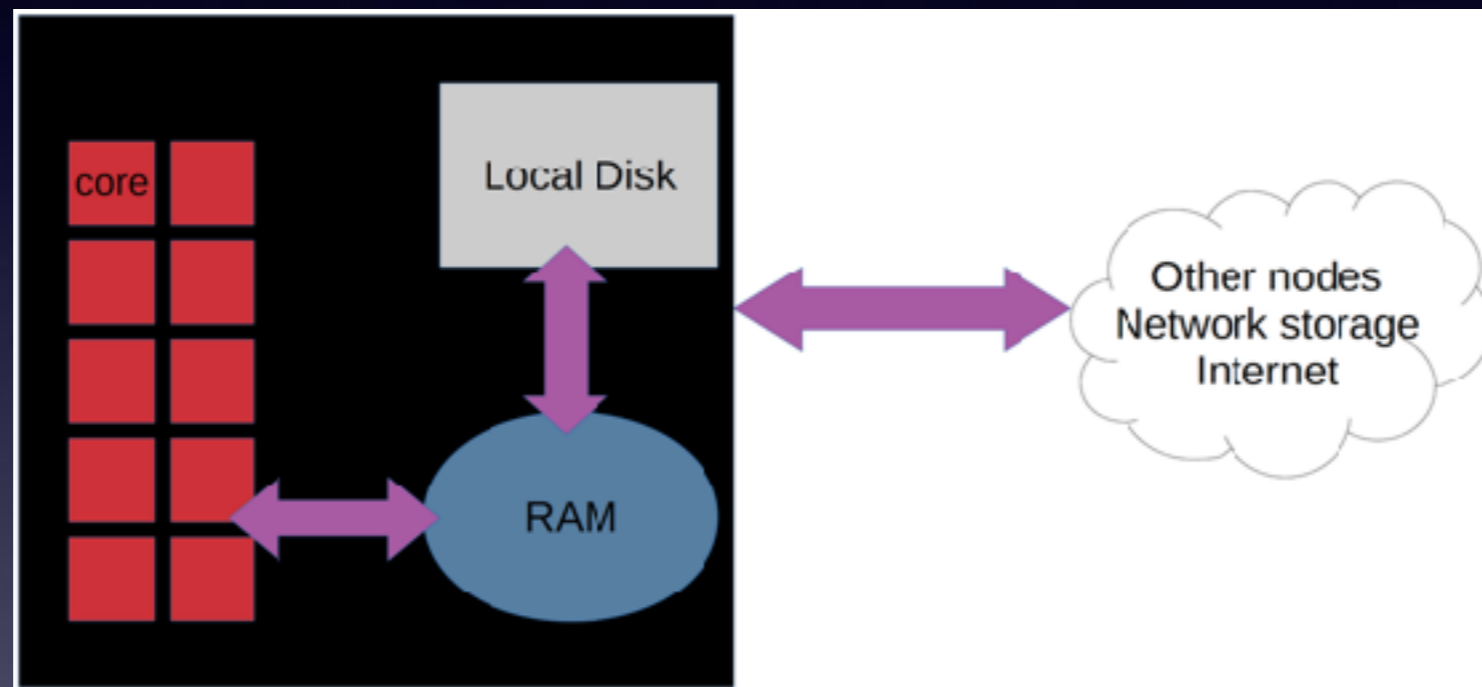
- “module spider <name-of-software>” to search
- Then “module load <name-of-software>” to activate a software environment

Break time

# Anatomy of a cluster



# Anatomy of a node



- CPU core — hardware that runs 1 thread
- Thread — 1 sequence of operations in software

# Running software

- Anything that takes more than a few minutes to run must be run in a **job**.
- A job is either **interactive** (you log directly into a compute node) or a **batch job script**.
- Jobs occupy a number of cores until the job is done, canceled, or times out.
  - Check if the job actually \*used\* the occupied cores using our **jobstats** utility (at UPPMAX)

# Slurm basics

- Slurm schedules jobs on resources
  - -A project
  - -p partition
  - -n #cores OR -N #nodes
  - -t maximum runtime



# Interactive jobs

- For direct login to compute nodes
- Limited length and size
- ```
$ interactive -A snic2021-22-606 -p  
core -n 1 -t 10:00
```

# Batch jobs

- Write “job.sh”
- Submit to queue with `sbatch`

```
#!/bin/bash
#SBATCH -A staff
#SBATCH -p core
#SBATCH -n 1
#SBATCH -J Template_script
#SBATCH -t 00:10:00

# load some modules
module load bioinfo-tools

# go to some directory
cd ~/uppmxintro/uppmx_tutorial

# do something
echo Hello world!
```



# NLP at UPPMAX

- We have a preinstalled Python ML module containing e.g. PyTorch, Tensorflow, scikit-learn and configured to work on CPU (Rackham) and GPU (Snowy)
- It is possible to run jobs interactively through Jupyter Notebooks in your browser. See Python HPC course for more info: <https://uppmax.github.io/HPC-python/index.html>
- If you prefer R, we also have that as well as Rstudio and most packages available on CRAN preinstalled

# Etiquette & Efficiency

- Don't bog down the log-in nodes with:
  - Excessively heavy or long-running programs
  - Excessive disk I/O
- Check job efficiency with “jobstats” tool and/or measuring runtime with fewer cores

# Disk usage guide

- At UPPMAX:
  - Backup system preserves old versions for 30 days
  - Put temporary and volatile files in a directory called “nobackup” — makes backup system better
- Use “pipes” and “named pipes” to avoid creating temporary files
- For jobs that do **intensive** file accesses, consider copying the files to **\$SNIC\_TMP**
- Always compress data
- Always delete unnecessary data
- Always move “finished” data sets off

Home

Scratch

Project

# Thank you for listening

Visit [uppmax.uu.se](http://uppmax.uu.se) for news and guides

Check [status.uppmax.uu.se](http://status.uppmax.uu.se) for system status