Portuguese Distributed Computing Infrastructure

INCD

Géant Agreement Experience

http://www.incd.pt

co-funded by:

Contract

- Select provider
  - Not all are active in Portugal
  - Difficult to navigate in the documentation
  - Difficult to compare the offers
  - Contract with AWS via TI Sparkle
- Bureaucracy
  - Sign Call Off Agreement + Order Form
  - Selected monthly payment → monthly

Conditions

- Price
  - 3% discount over AWS street price
  - On top of that 3% of administrative costs
  - Full data egress waiver

Service

- Support
  - Via TI Sparkle
  - Responsive
- Accounting and management
  - Not good cannot control usage via AWS tools
  - Need to use TI Sparkle tools to manage accounts
  - Need to use TI Sparkle tools to check consumption
  - These tools are not as good or complete as AWS

Troubles

- Minor troubles
  - TI Sparkle asked for some additional data
    - we complaint they accepted
  - Needed to sign GDPR amendment
    - when we initially signed GDPR was already into force
  - Invoices didn’t appear and then got complaints
    - had to ask for the invoices to be sent

AWS comparison GROMACS

performance and price comparison

GROMACS

- GROningen MAchine for Chemical Simulations (GROMACS)
- Originally developed in the Biophysical Chemistry department of University of Groningen
- Use GROMACS for simulations of proteins, lipids, and nucleic acids.

GROMACS
FAST. FLEXIBLE. FREE.

Motivation

- GROMAC’s is widely used by scientists
- Performance and Price comparicing:

  AWS EC2’s using Batch

  vs

  INCD physical batch server
- Our intention was using AWS lambda but resource limits on AWS lambda are now much stringent

Running with containers

- Used Docker container to encapsulate and execute GROMACS
- “Rootless containers with udocker” - Samuel Bernardo (17:30h)

UDOCKER

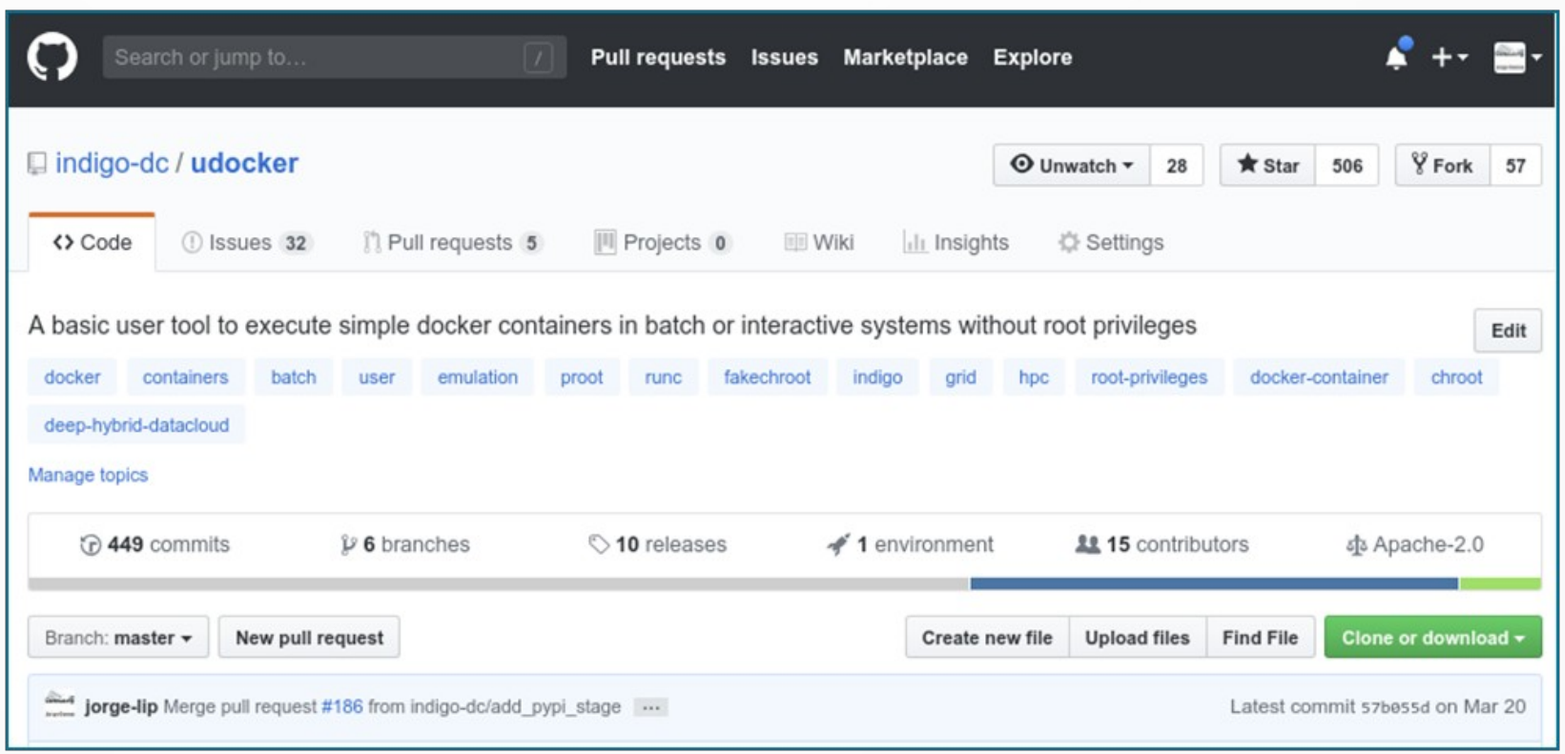

https://github.com/indigo-dc/udocker

udocker on farm

- “basic user tool to execute simple docker containers in user space without requiring root privileges”

- 8 physical cores per each container – 100% cpu usage – 100 runs

(8 containers / each run)

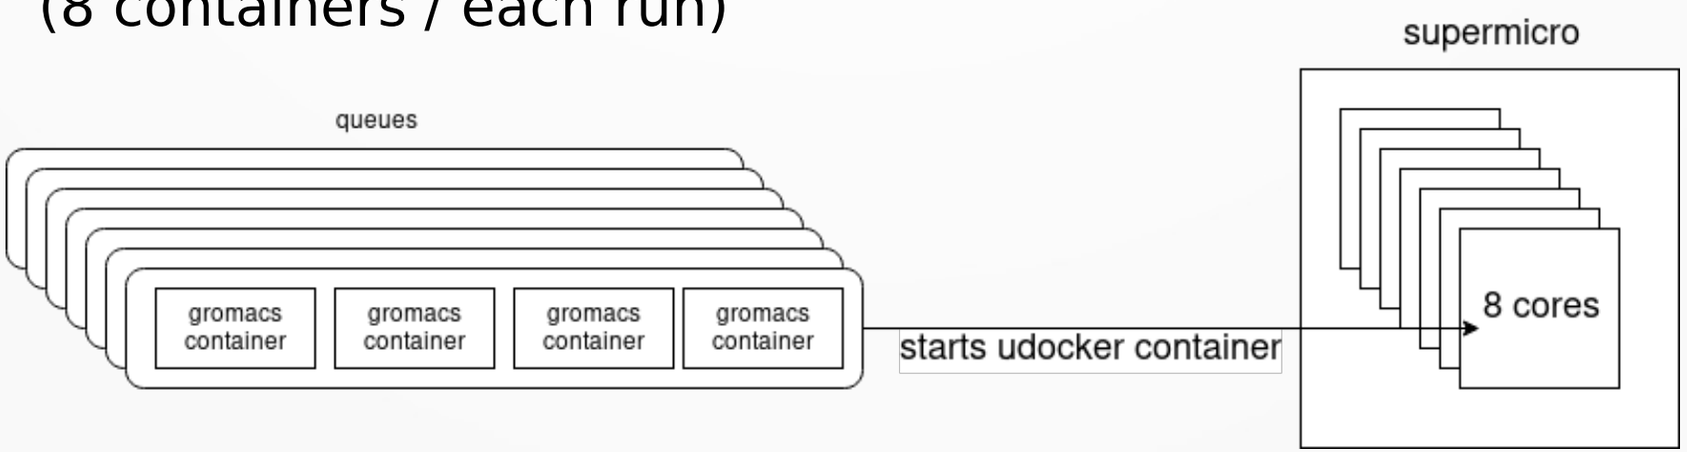

scar on aws

- “framework to transparently execute containers out of Docker images in AWS Lambda and AWS batch using udocker”

- One **c4.2xlarge** per each container – 100% cpu usage – 100 runs

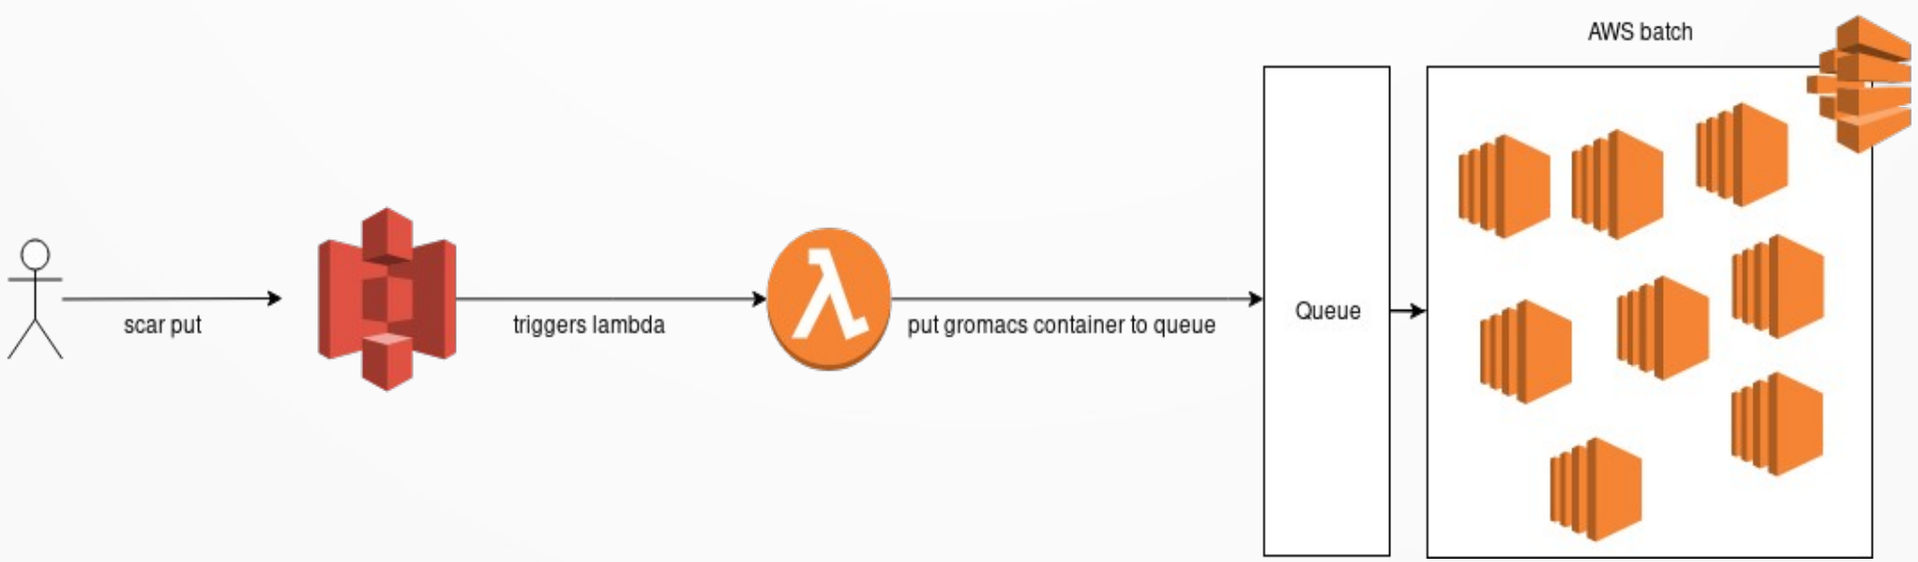

GROMACS performance

Average:

Farm:

1h17m19s

AWS:

1h02m44s

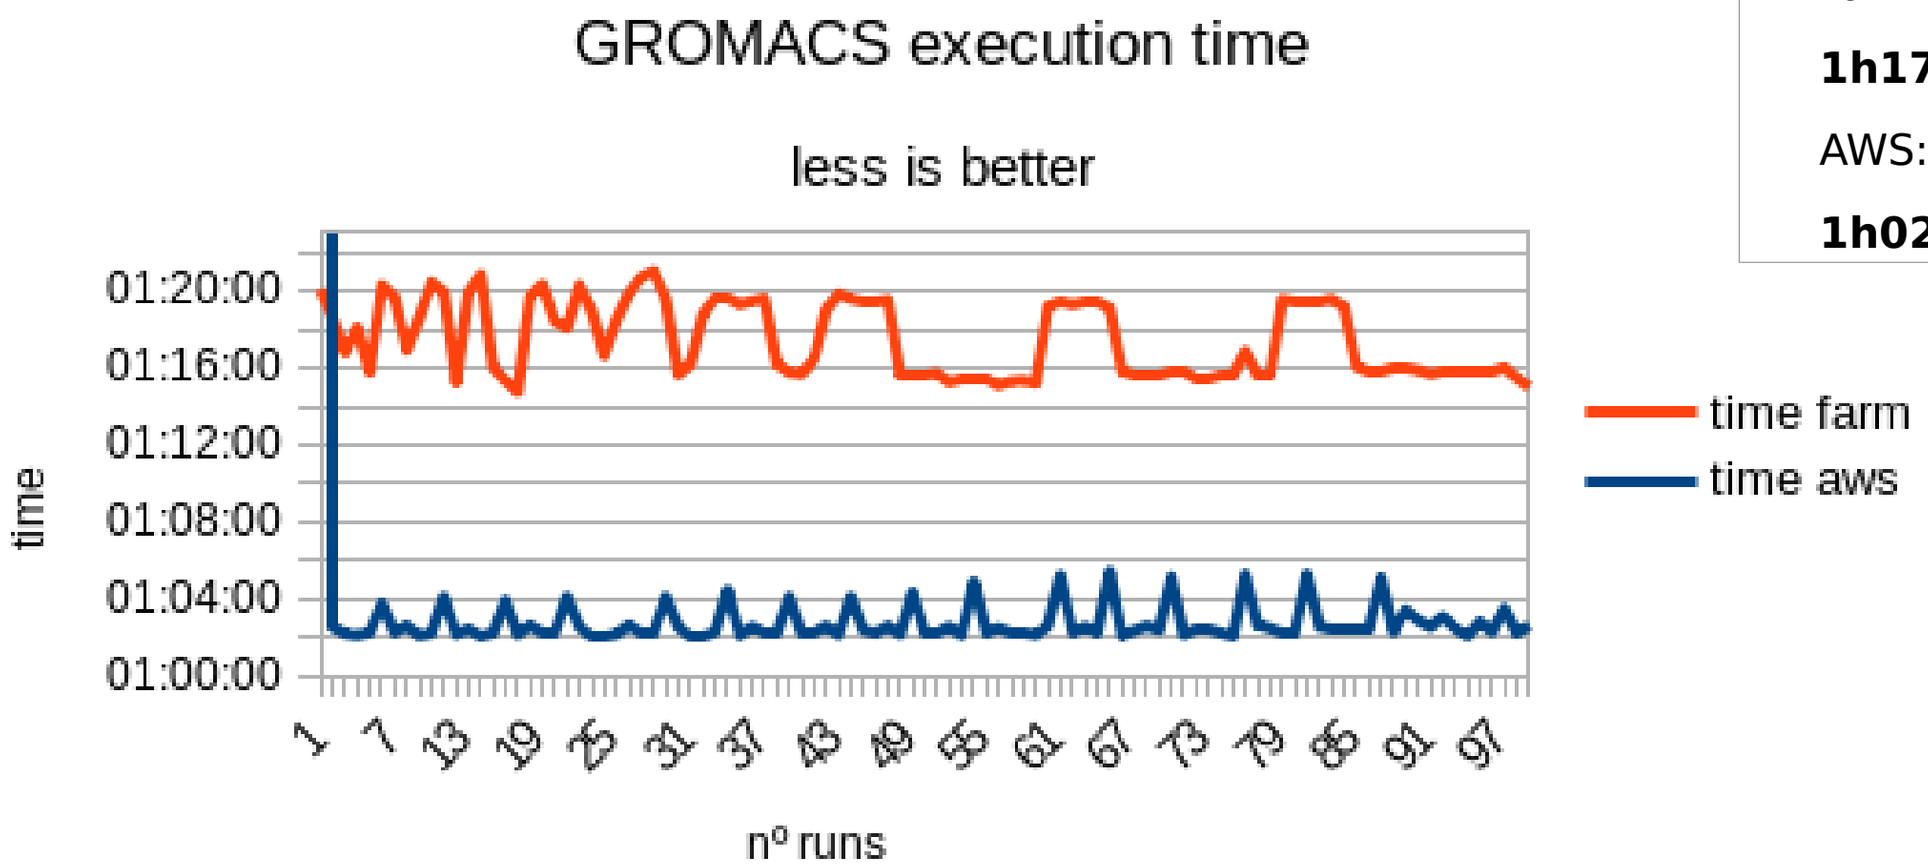

GROMACS performance

Average:

Farm:

18.58ns/day

AWS:

22.95ns/day

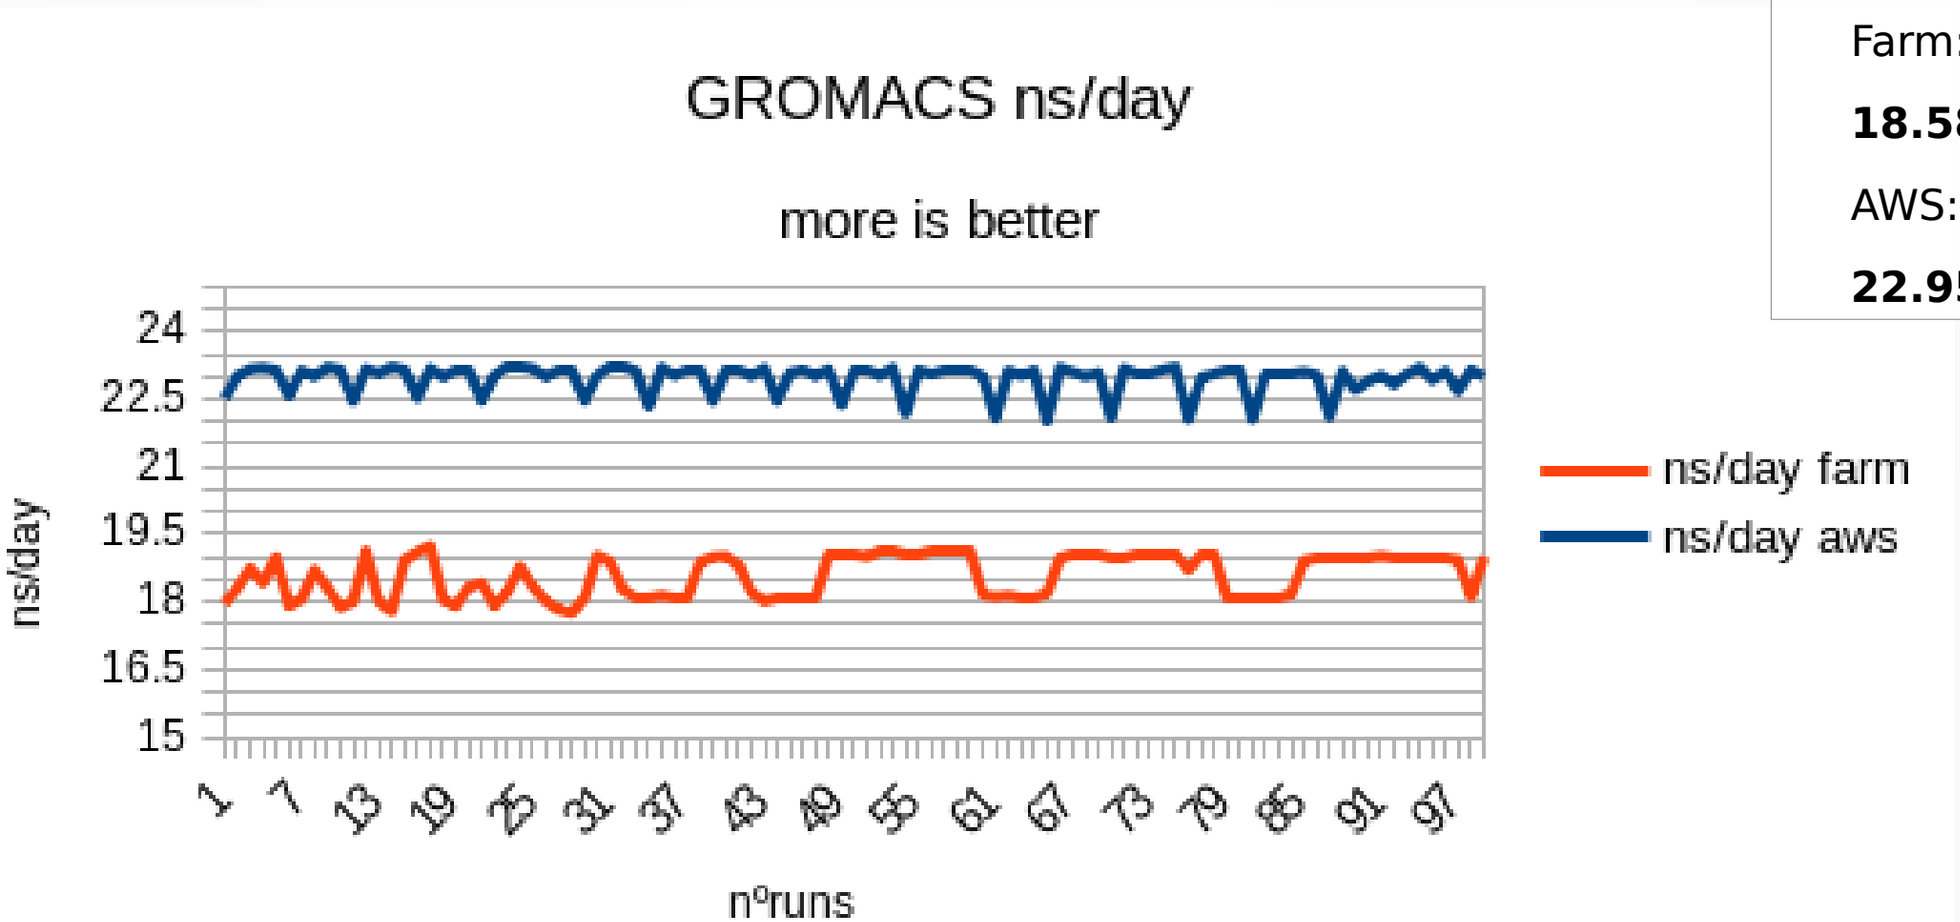

CPU's

AWS

- **Intel Xeon E5-2666 v3**

  Clockspeed: 2.9 GHz

  Turbo Speed: 3.3 GHz

  No of Cores: 8

  Memory: 15GB for 8 Cores

  Used by gromacs: 8 cores

INCD

- **AMD EPYC™ 7501**

  Clockspeed: 2.0 GHz

  Turbo Speed: 3.0 GHz

  No of Cores: 64

  Memory: 48GB for 8 Cores

  Used by gromacs: 8 cores (8/64)

AWS Price

- Price of 8x c4.2xlarge EC2 instances, with 100% CPU utilization for 1 month:

  around **2330 €**

- AWS support on gromacs:
  - None

- What’s included:
  - EC2, Batch, CloudWatch

- Required:
  - Some AWS knowledge

INCD Cost

- Price of 1 Supermicro 1123US-TR4, with 100% CPU utilization for 1 month:

Energy:

Comsumption per hour 480 W

In 24 hours:
480 W * 24h = 11.52 kWh

Price per KWh:
0.132€ * 11.52 kWh = 1.52€

Price of the machine in one month:
1.52€ * 30 days = ~ **46€**

% of other energy costs: **~ 6€**

PUE: **1.5**

Other:

The machine + network + other hw:
~ 15000€

Amortization of 5 years (60 months).

Amortization by month:
15000 / 60 = ~ **250€**

TOTAL = (46€ + 6€) * 1.5 + 250€ = 328€
NORMALIZE TO AWS = 328 * 1.2 = 394€

INCD Cost

- Price of 1 Supermicro 1123US-TR4, with 100% CPU utilization for 1 month:

  around 394**€**

- INCD support on gromacs:
  - By e-mail, phone and slack

- What's included:
  - 8h/day support, application support

- Required:
  - Basic CLI knowledge

Final thoughts

- Comparison for 1 month of Gromacs 64 CPU COREs
  - AWS: 2330€
  - INCD: 394€
- AWS advantage
  - Very large capacity on demand
- INCD advantage
  - Requires less skills, application, support, cost

Recommendations

- Consider all aspects including costs and benefits carefully
  - Ocasional or permanent use
  - Development/adaptation of applications
  - Technical requirements and resource consumption
  - Privacy, security, applicable law
  - Confidentiality and data value
  - Reliability and trustworthiness of the infrastructure
  - Competences, training and knowledge in the organization
  - Cash flow and funding to pay services
  - Evaluate risks
  - Evaluate the cost

Thank you

AWS dedicated instances

- Price of 8x c4.2xlarge EC2 dedicated instances:
  - $2567 all upfront 1 year
  -
  - 8 x (2567 x 0.90 x 1.23) = 22733 Eur
  - 22733 / 12 = 1894 Eur / month

AWS spot instances

- Price of 8x c4.2xlarge EC2 spot instances:
  - $0.1234 per hour
  - 8 x (0.1234*24*30 x 0.90 x 1.23) = 787 Eur / month

AWS c5n.18xlarge

- Price
  - $3.888 per hour
  - 8 x (3.888*24*30 x 0.90 x 1.23) = 3098 Eur / month

INCD other costs

- Human resources
  - You pay both on AWS and on INCD
    - AWS: your effort of managing VMs, CloudWatch, Batch, Orchestration, etc
    - INCD: our effort of managing the infrastructure + user batch submission
  - INCD human resources in Lisbon center
    - 5x FTE / nodes =~ 120€ / month / node
    - The larger the national infrastructure more efficient it will be