HPC Infrastructure for Energy: Past, Present, and Beyond

HPCKP Annual Meeting 23

Felipe Portella
IT Advisor, Petrobras
Research Visitor, BSC
PhD Candidate, UPC

May/2023

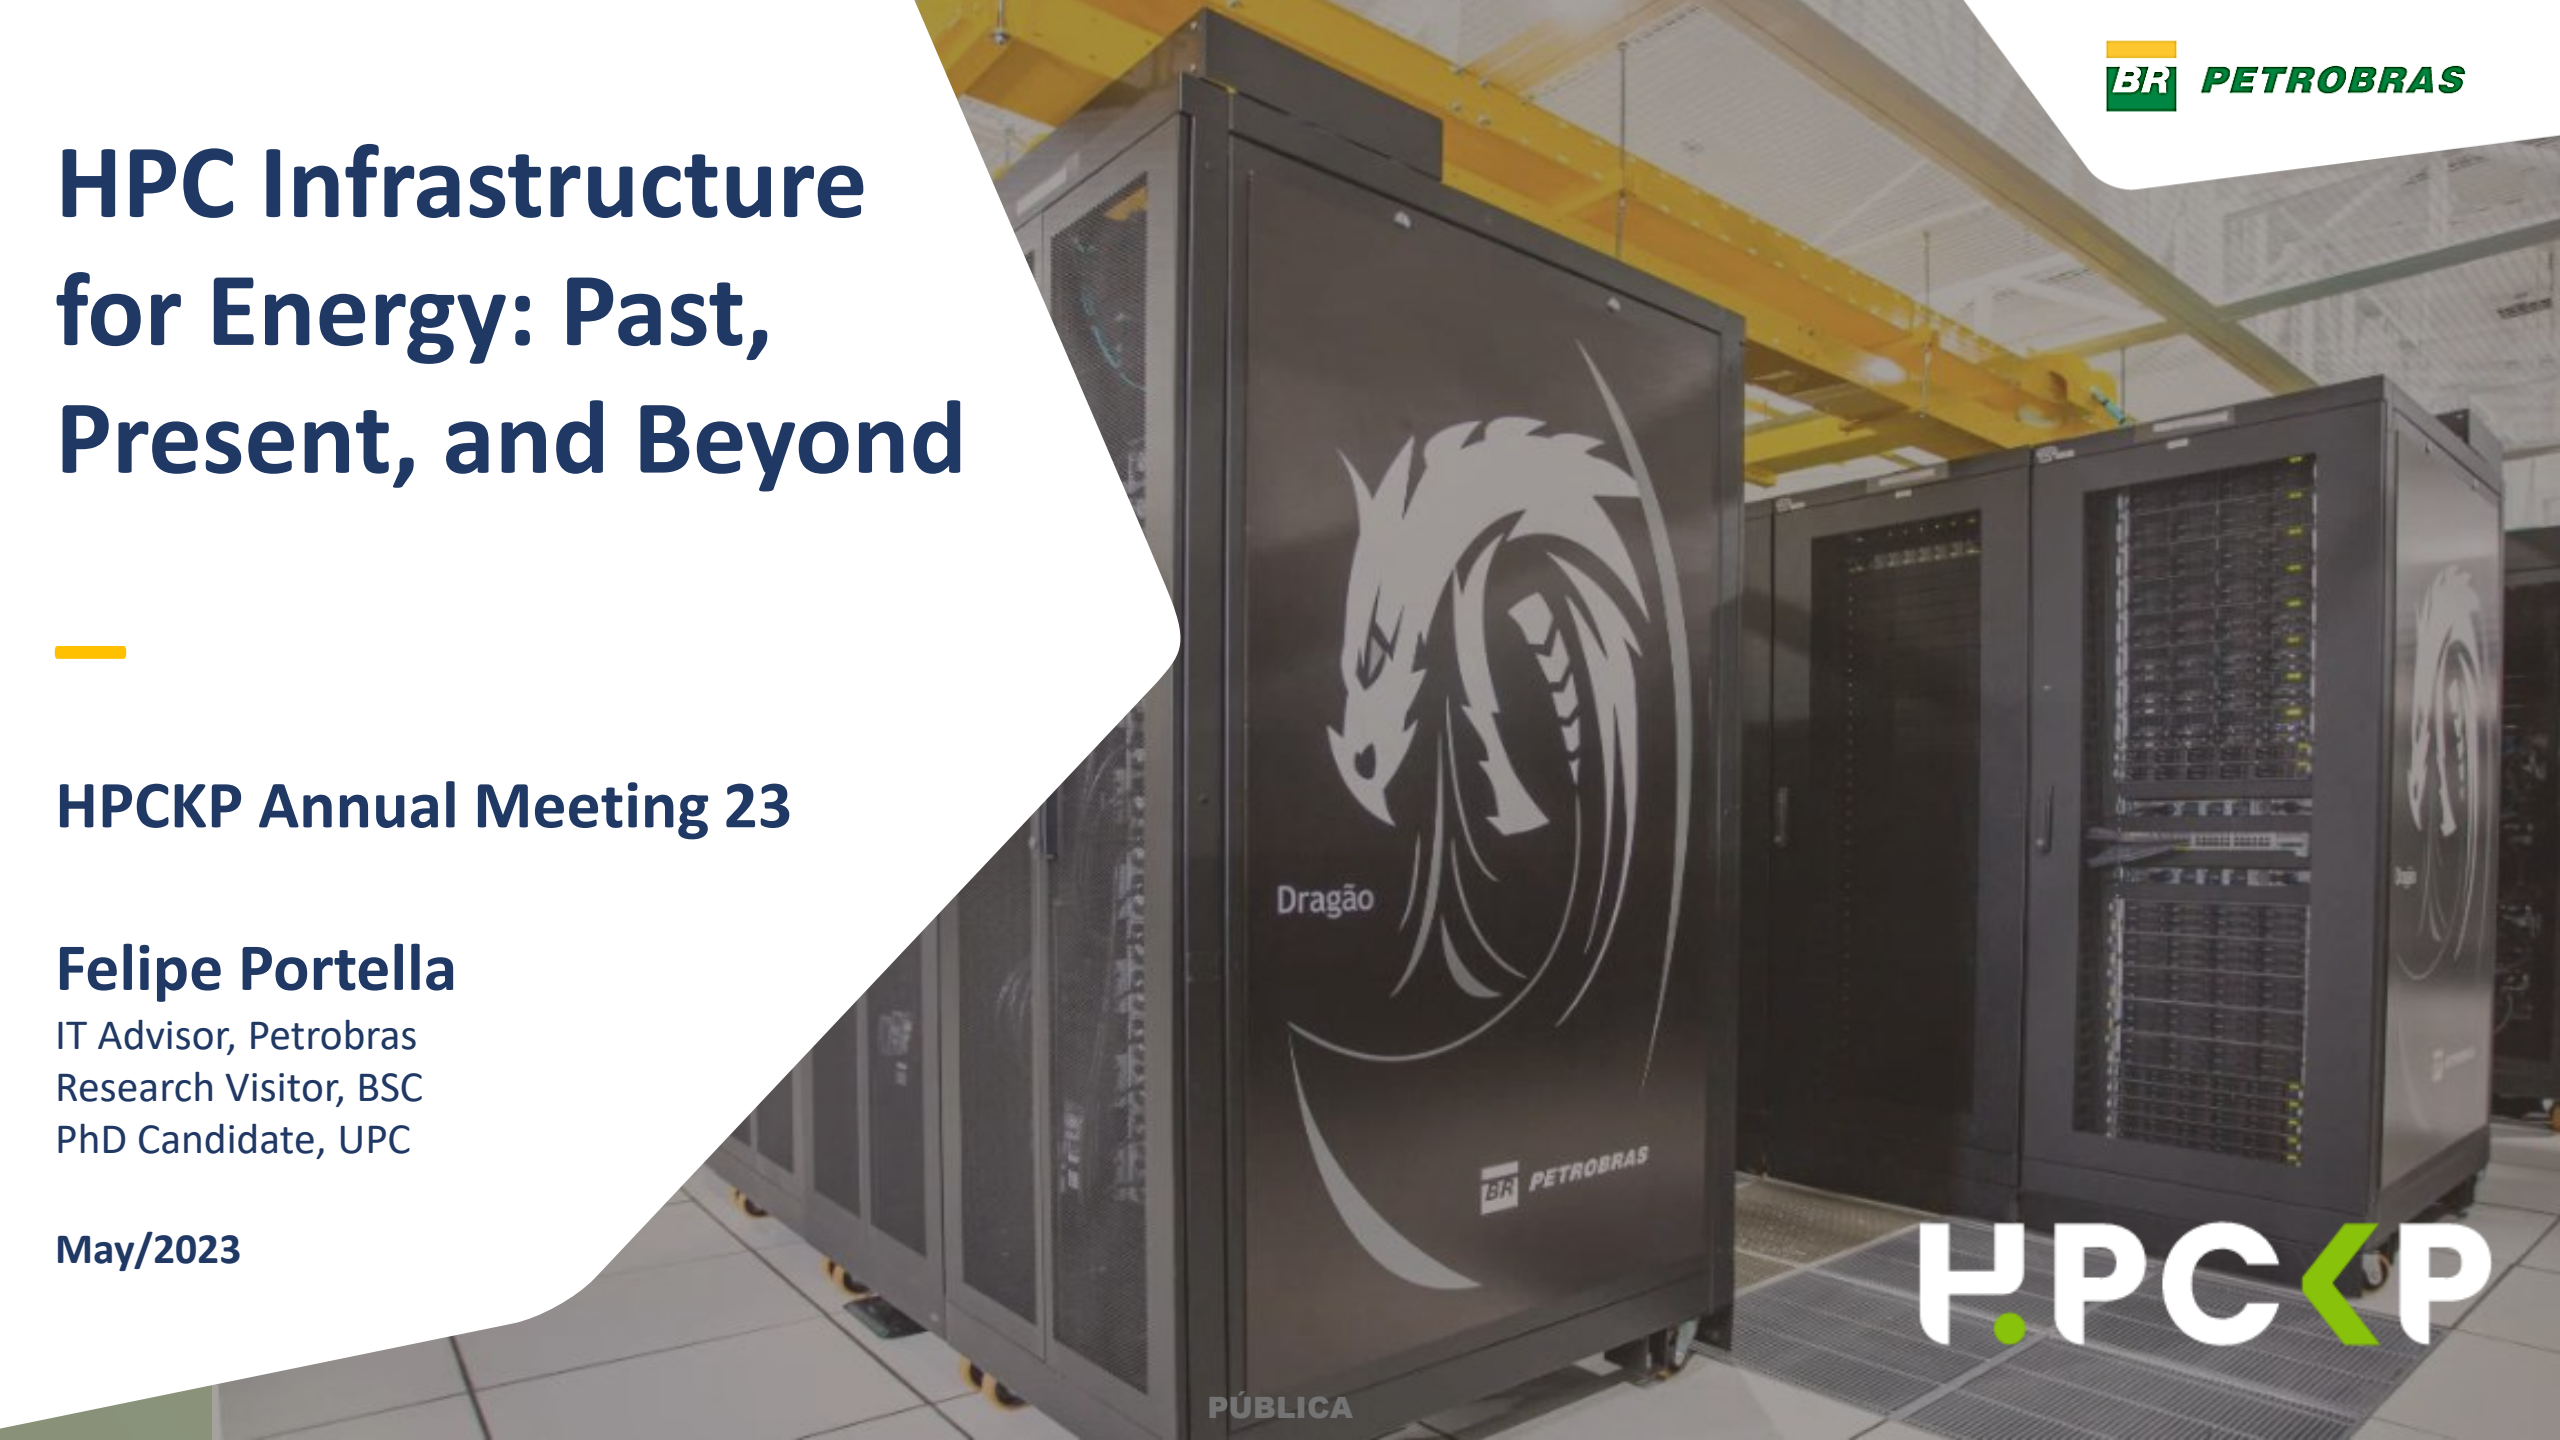

Agenda

- Why we use HPC at Petrobras
- How HPC evolved at Petrobras
- Our current HPC infrastructure
- Our future trends
- Conclusions

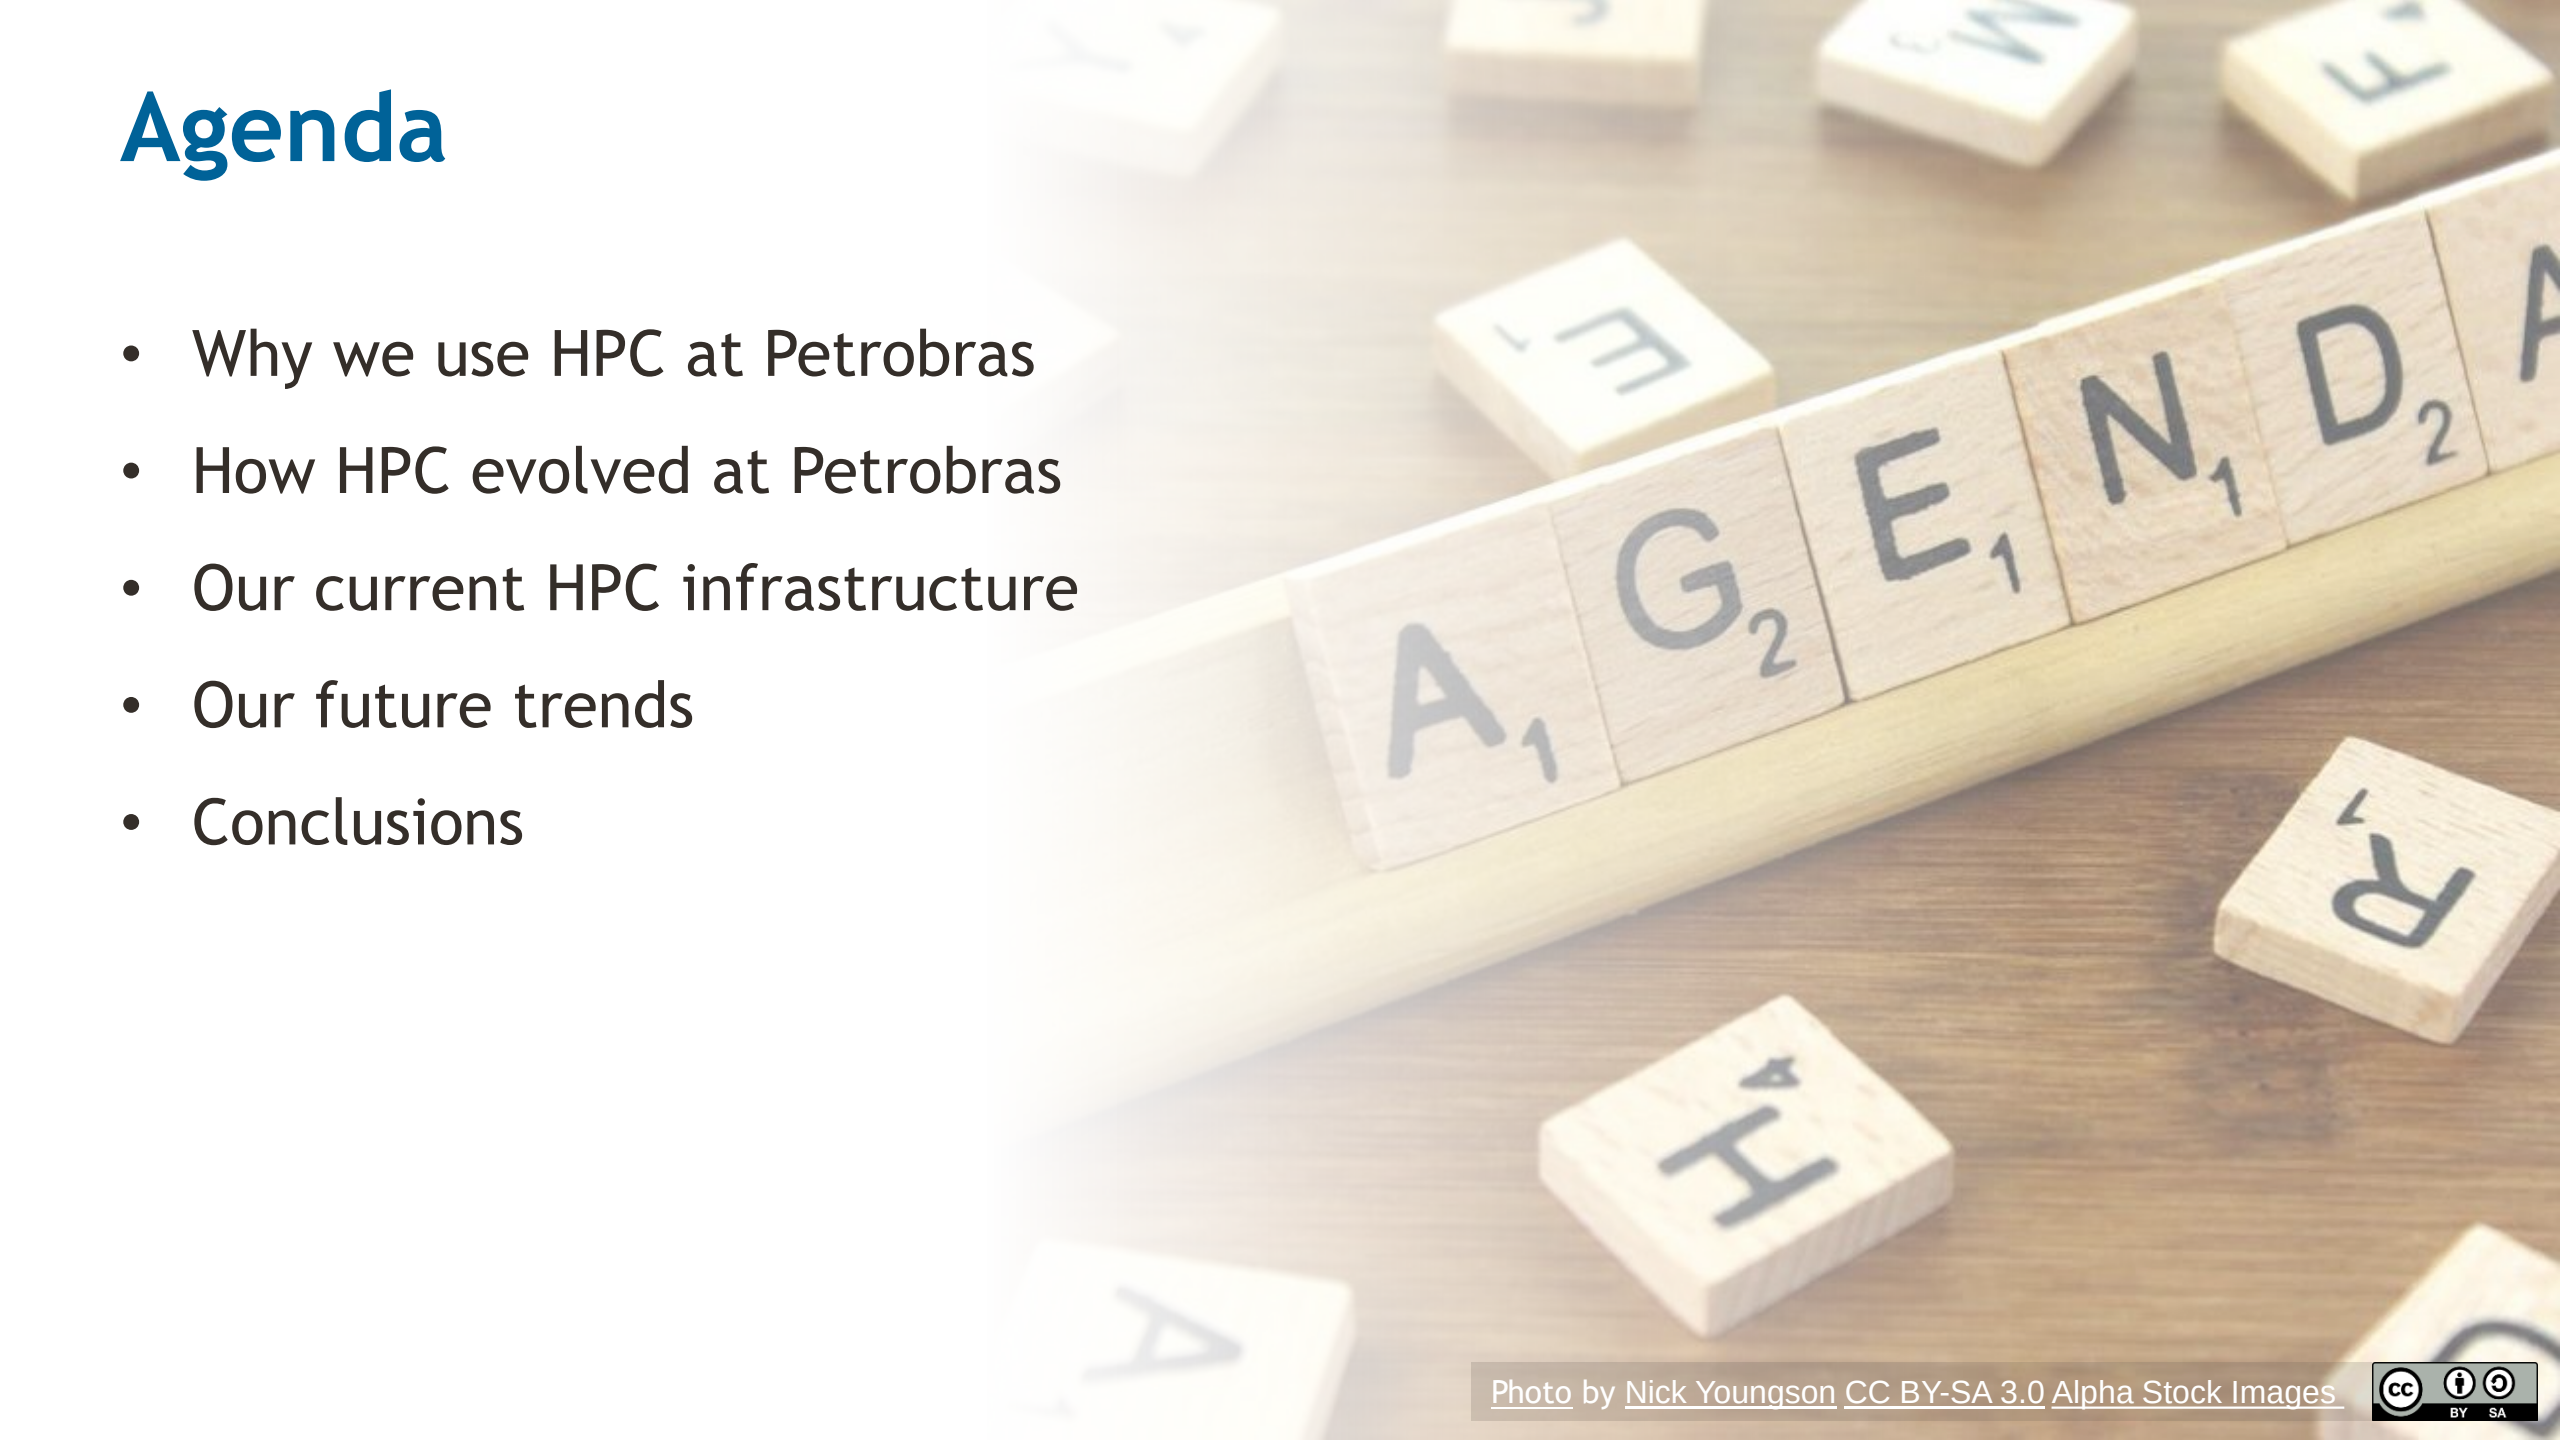

Agenda

- Why we use HPC at Petrobras
- How HPC evolved at Petrobras
- Our current HPC infrastructure
- Our future trends
- Conclusions

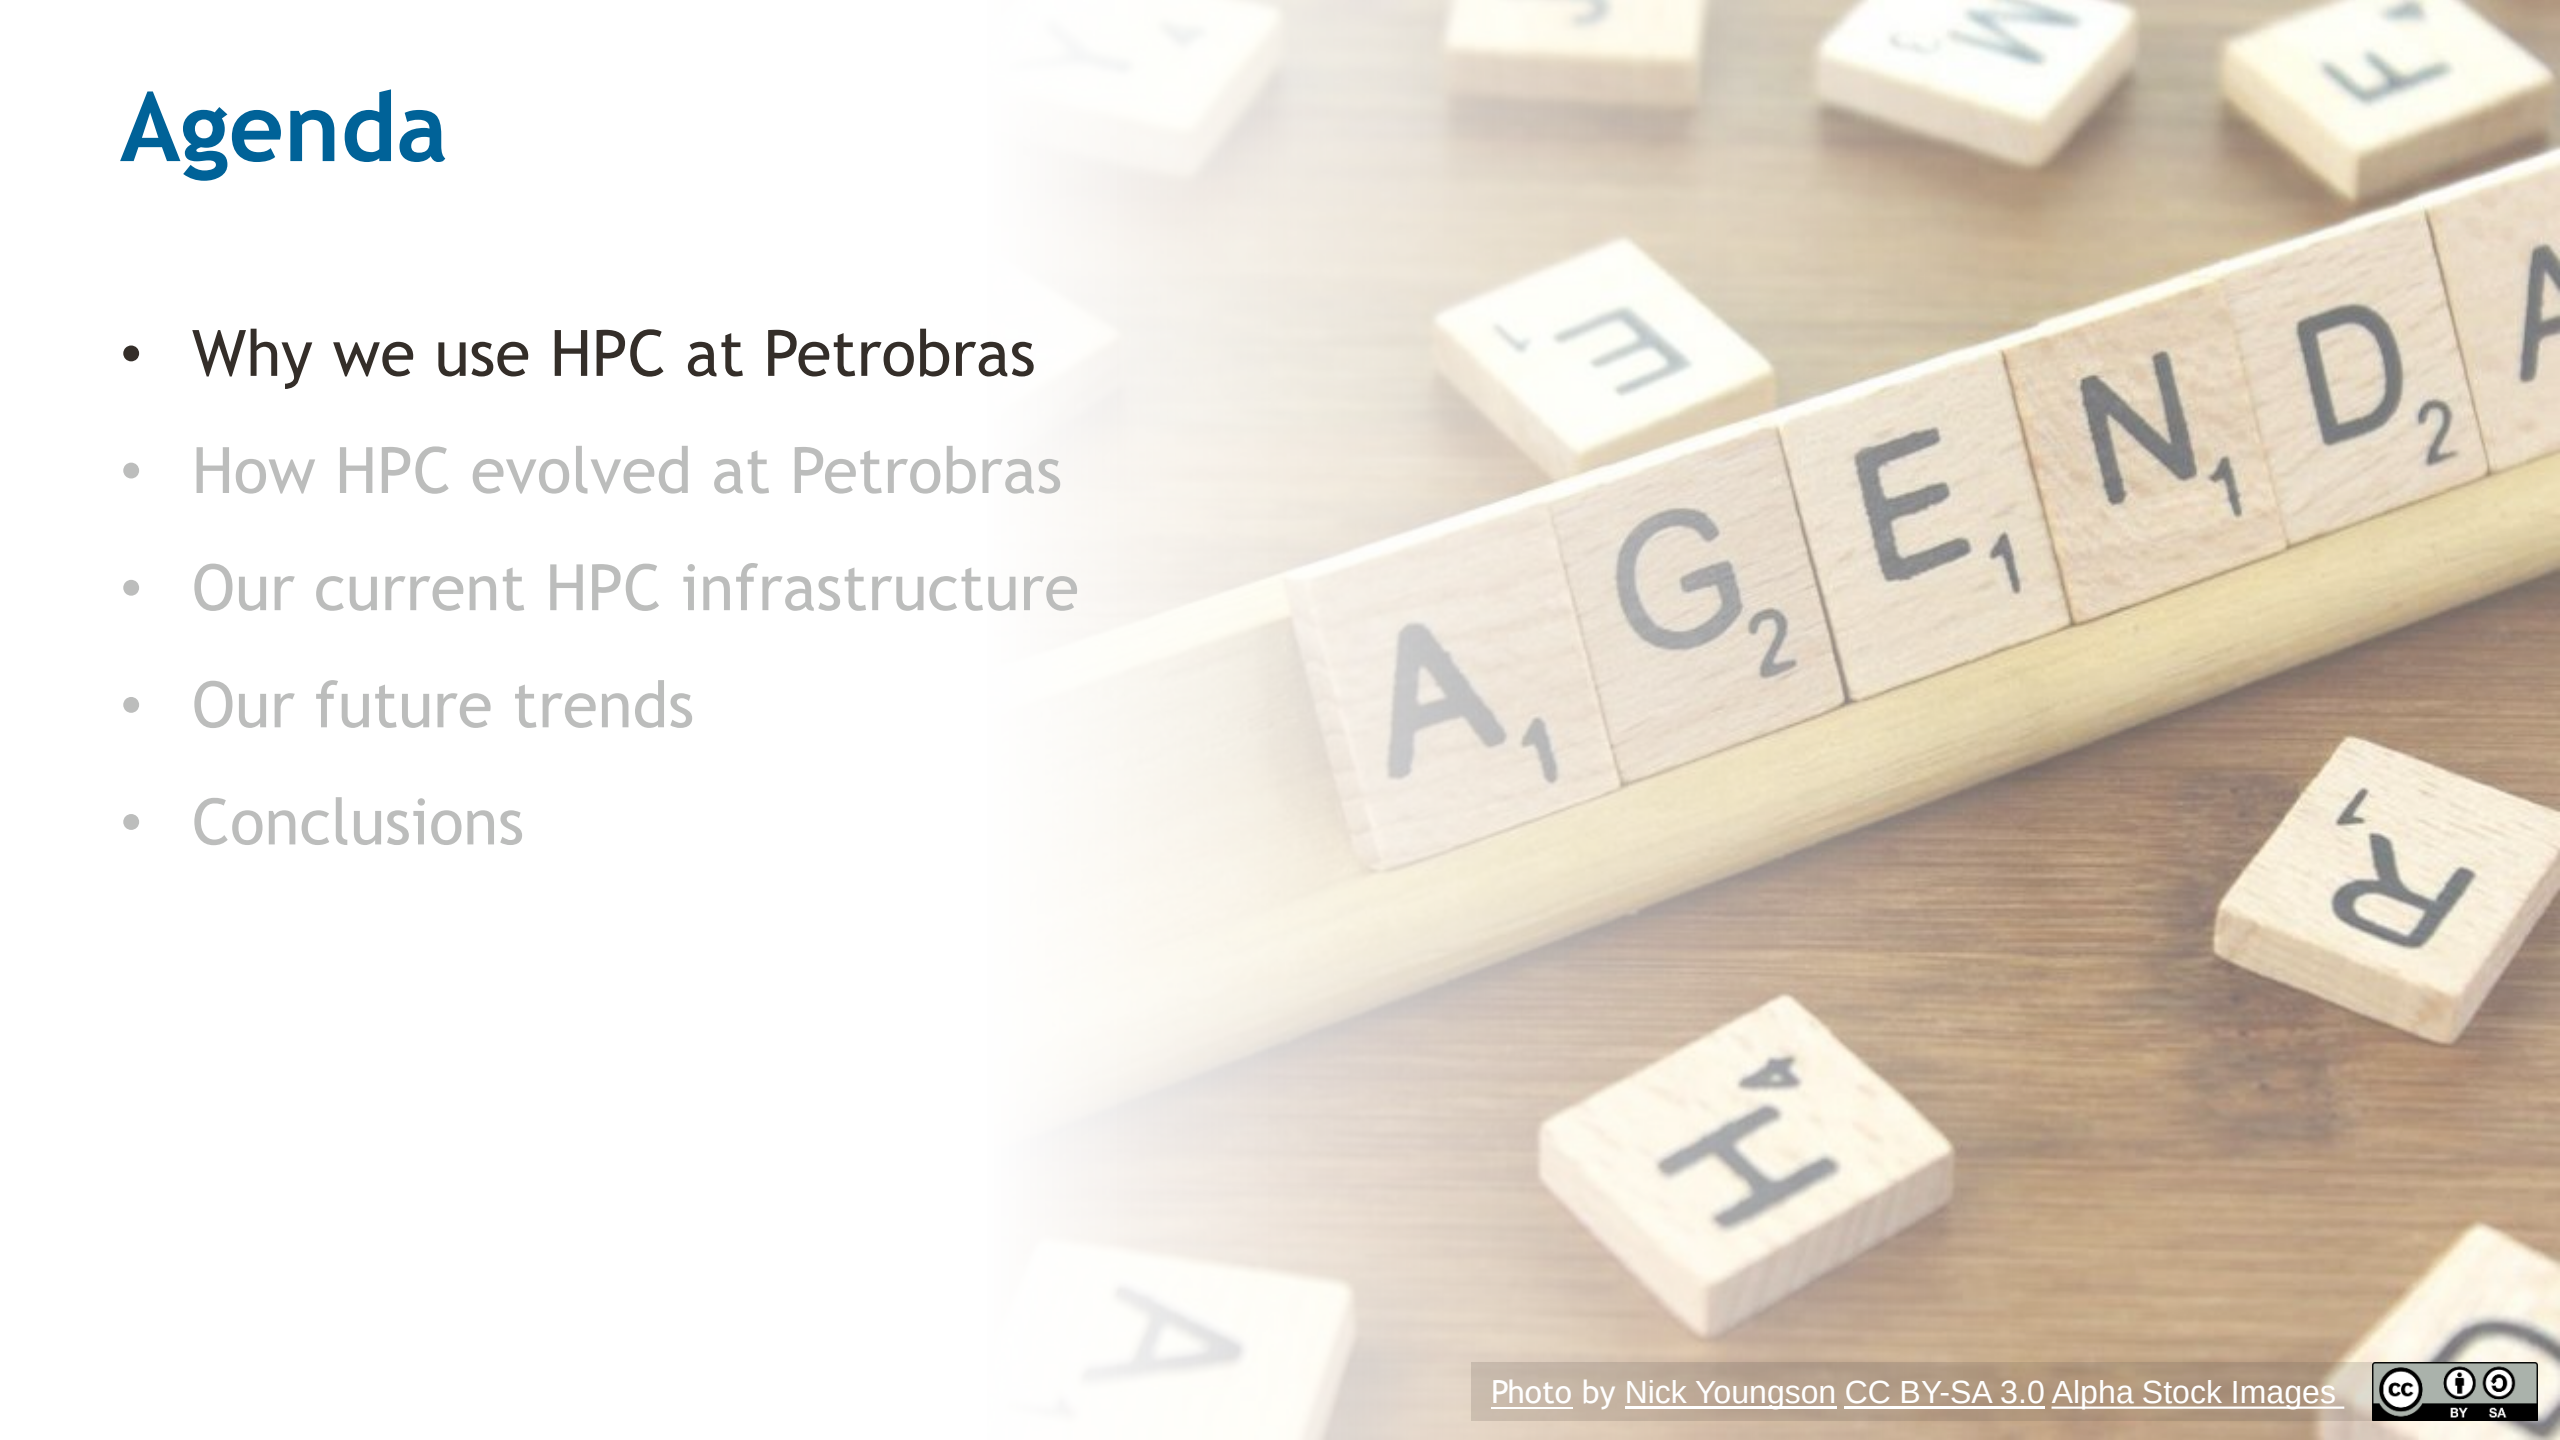

PETRÓLEO BRASILEIRO S.A. Quick Facts

- Brazilian Energy Company
- Oil and Gas Exploration and Production as core business
- Reference in ultra-deepwater exploration
- Established in 1953
- More than 45,000 employees in 2021
- Among the top10 energy companies in the world [1] and 128th on Fortune Global 500 2022 [2] in revenue

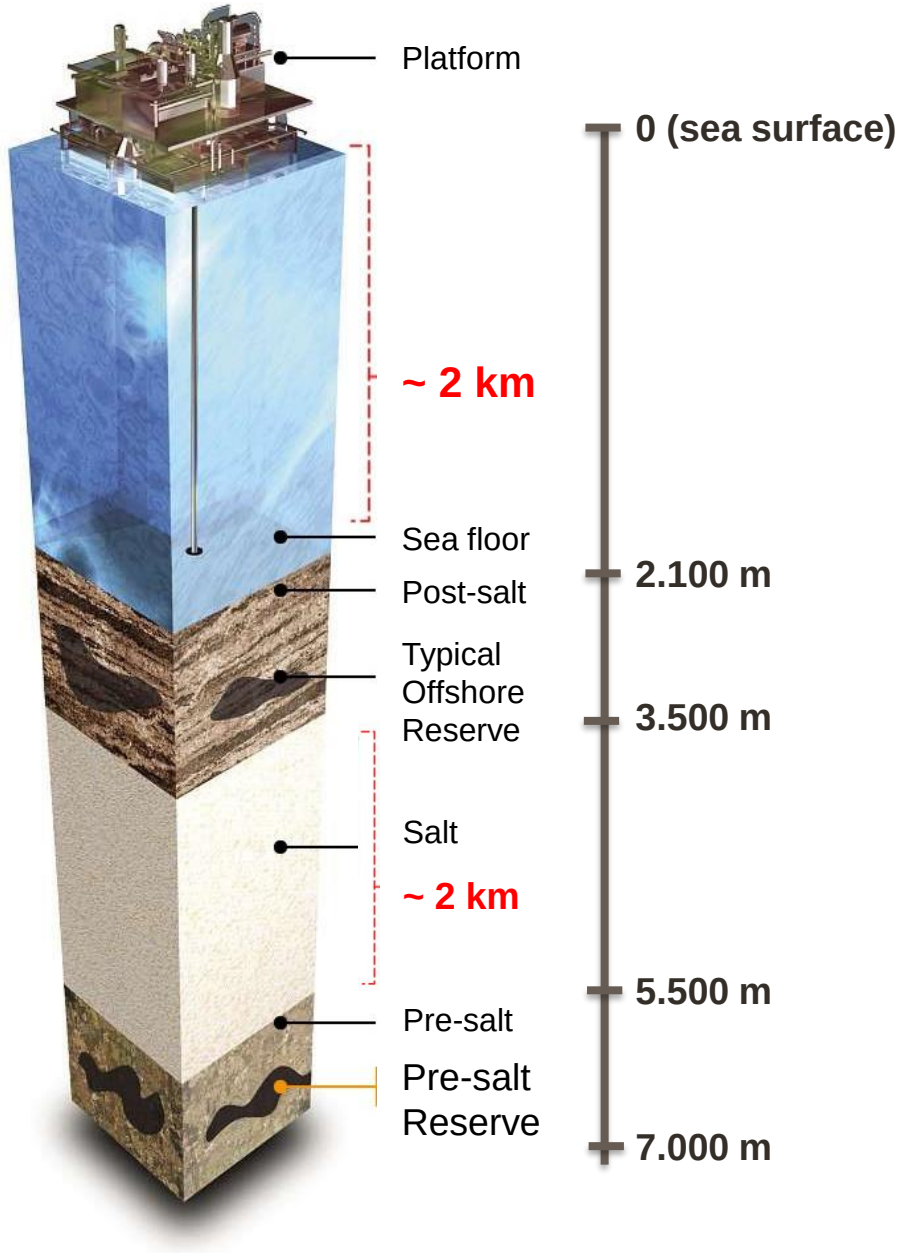

1. http://www.value.today/world-top-companies/energy
2. https://fortune.com/global500

Why HPC at Petrobras?

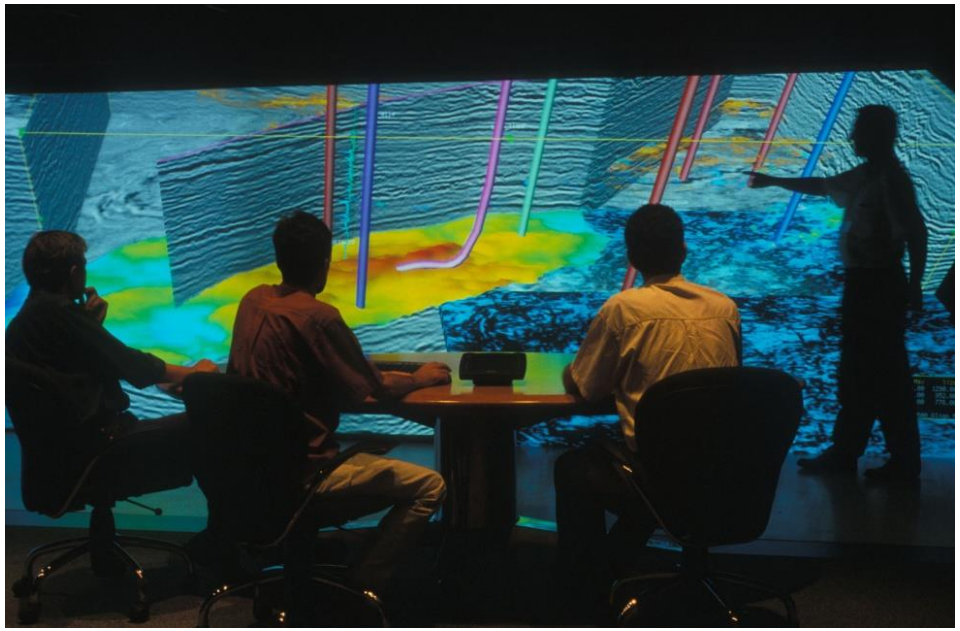

Seismic Processing

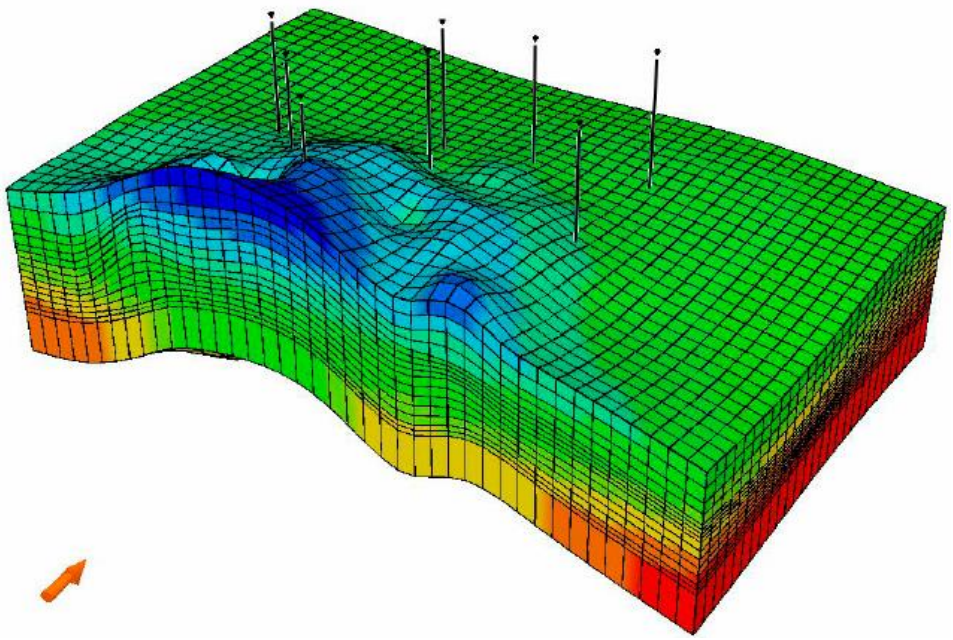

Reservoir Simulation

Each deep-water well typically costs between US$70 and US$100 million (1)

Petrobras will drill more than 300 offshore wells in the next 5 years (2)

(1) https://www.agenciapetrobras.com.br/pt/inovacao/petrobras-monta-supercomputador-para-desenvolver-tecnologias-18-01-2023/
(2) Petrobras Strategic Plan 2023-2027

Seismic Processing

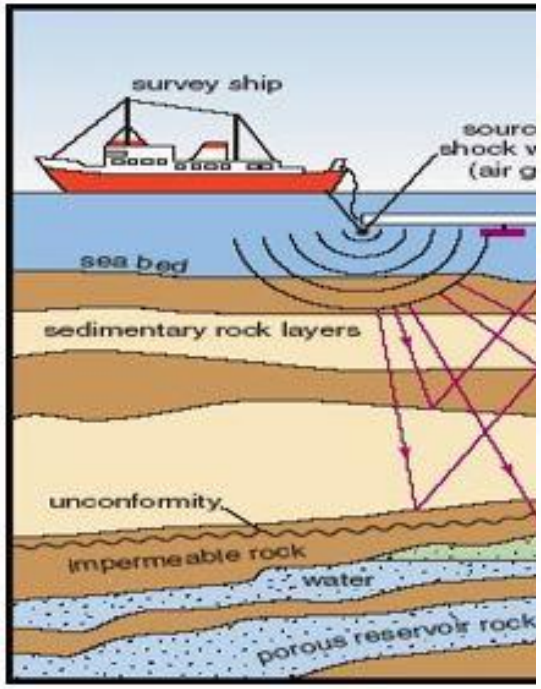

Credit: https://indogeologist.blogspot.com/

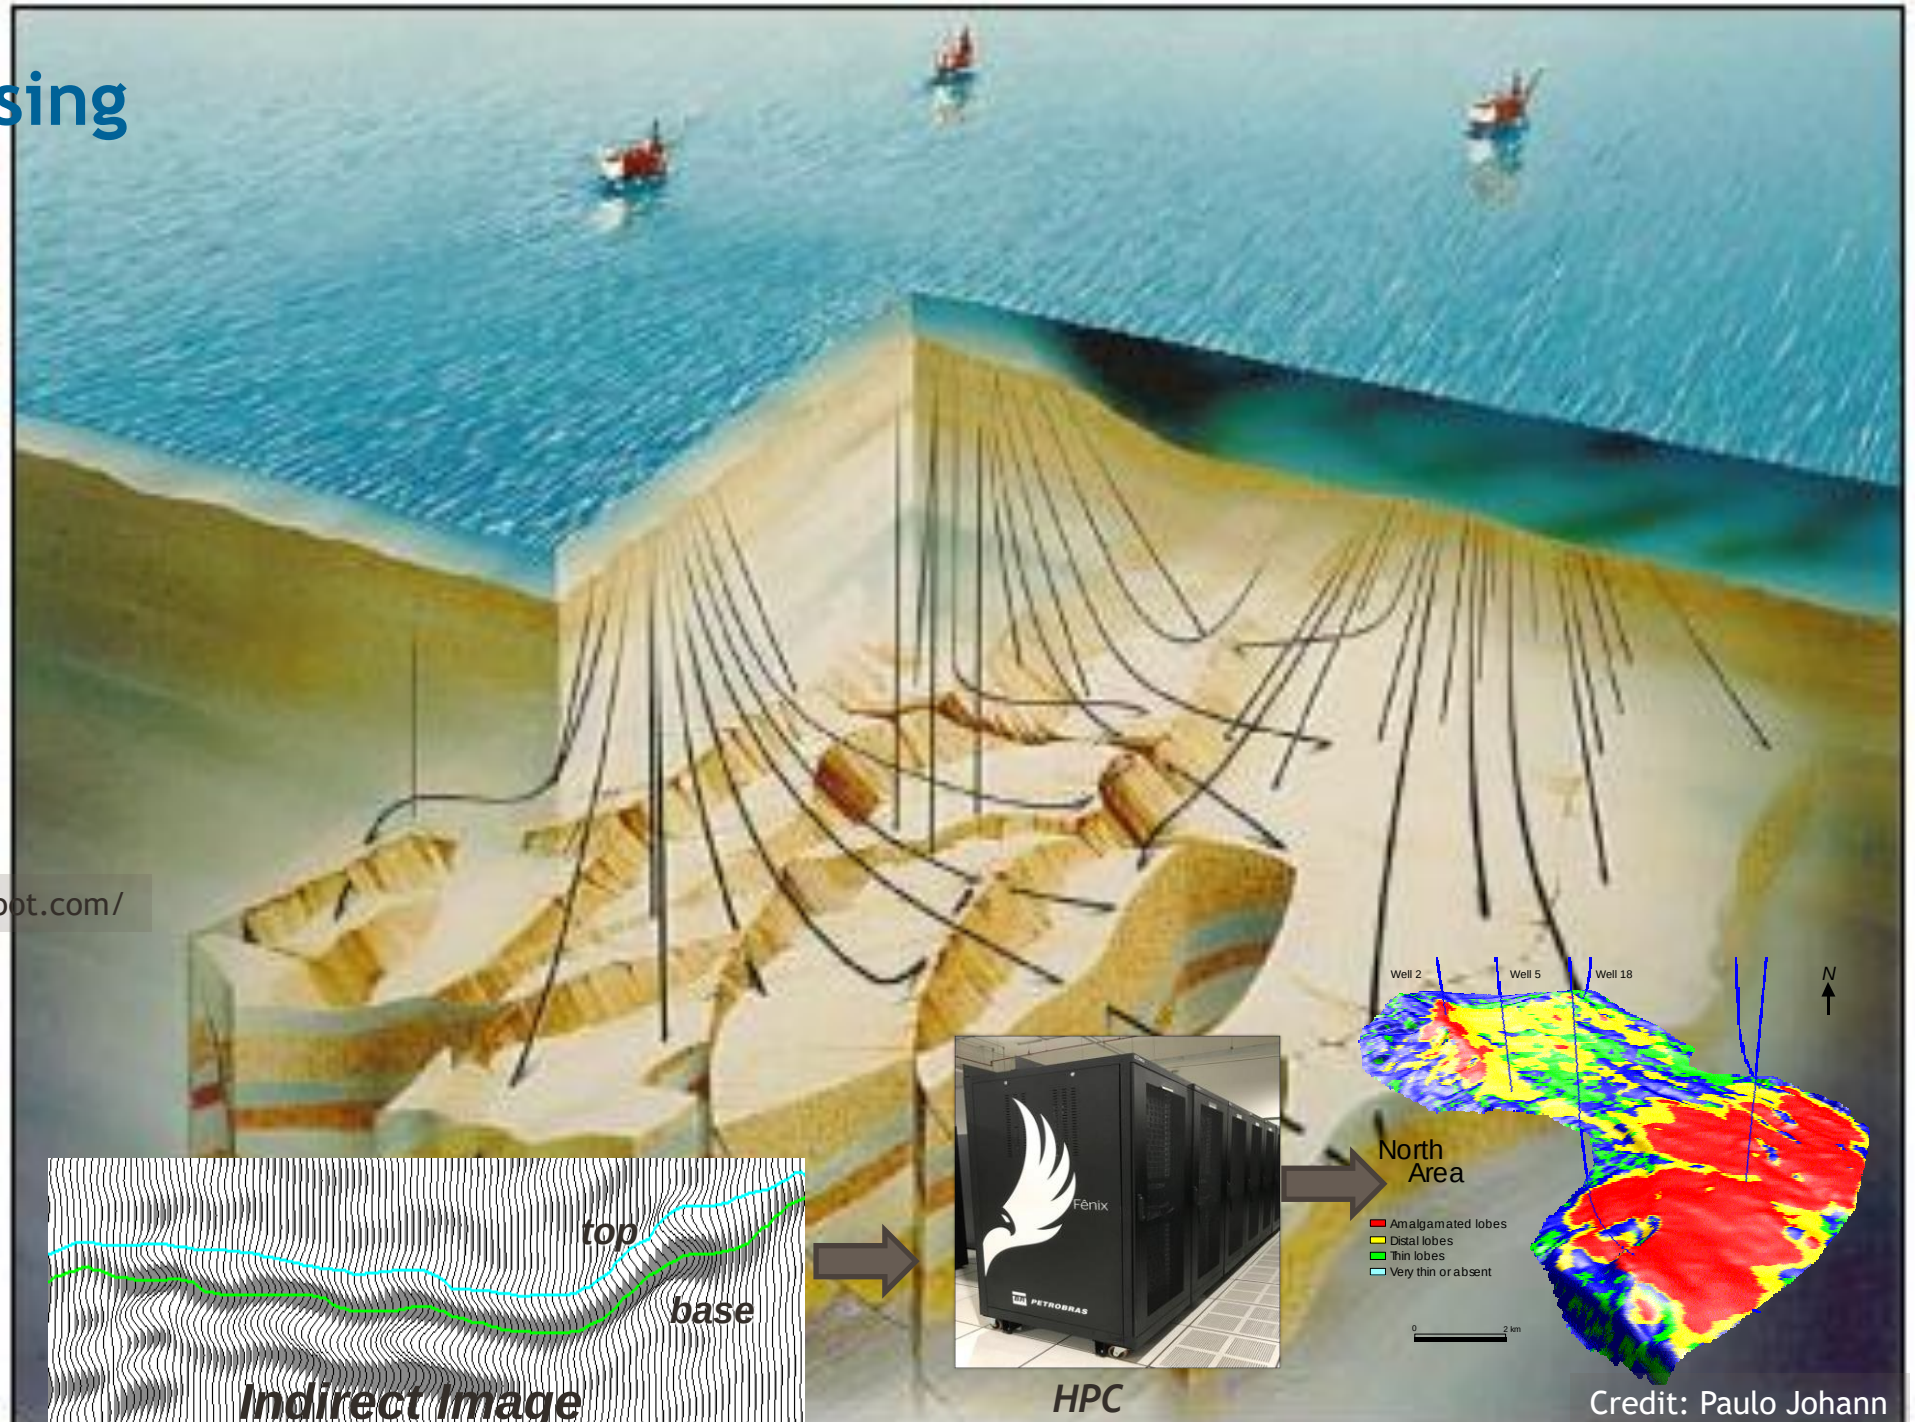

Credit: Paulo Johann

Reservoir Engineer Challenges

Exploitation Strategy

- **Wells**
  - How many?
  - Where should they be placed?
  - What type (vertical, horizontal, ...)?
  - Etc.
- **Recovery methods**
  - Primary depletion?
  - Injection of water or gas?
  - How to control the wells?
  - Etc.
- **Forecast behavior**
  - What are the final recovery volumes?
  - How about the cash flow?
  - Etc.
- **Etc.**

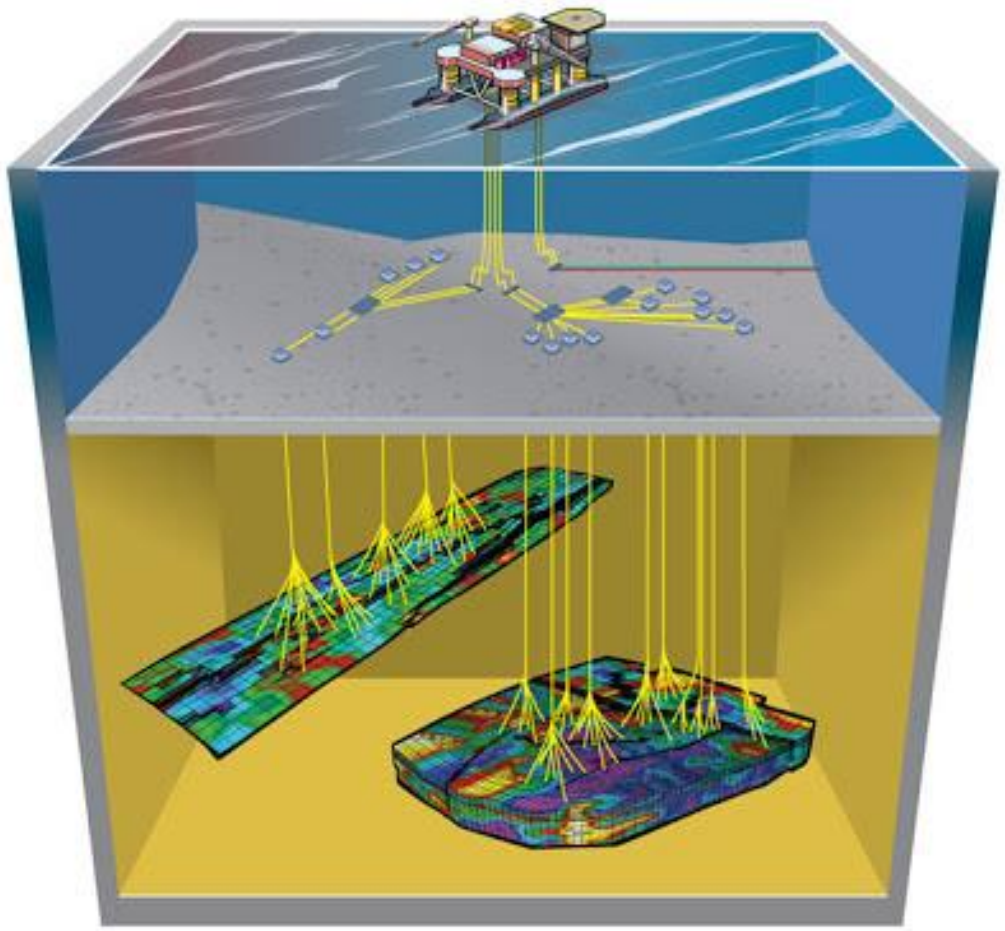

Image Source: PetroBlogger.com

Our main HPC workloads

PETROBRAS

- Seismic
  - Single precision (FP32)
  - Very very big datasets
  - Mix of CPU and GPU Jobs (most Processing capacity is GPU)
  - GPU jobs (typically) with Infiniband
  - Most important compute intensive imaging algorithms use GPUs and are developed in-house
  - Each job takes weeks
  - Some GPU jobs may use hundreds of GPUs (even thousands…)
- Reservoir
  - Double precision (FP64)
  - Not so big datasets
  - No GPU (there are exceptions)
  - Typically, commercial apps
  - Shared Memory (little need for high speed interconnect – with some exceptions)
  - Each job takes hours
  - Each workflow runs batches of dozens or hundreds of simultaneous Jobs
- Research, Petrophysics, Geomechanics, Multiphysics, Machine Learning, etc.

Agenda

- Why we use HPC at Petrobras
- How HPC evolved at Petrobras
- Our current HPC infrastructure
- Our future trends
- Conclusions

Once upon a time ...

Petrobras “HPC” in 1969: IBM 360-44

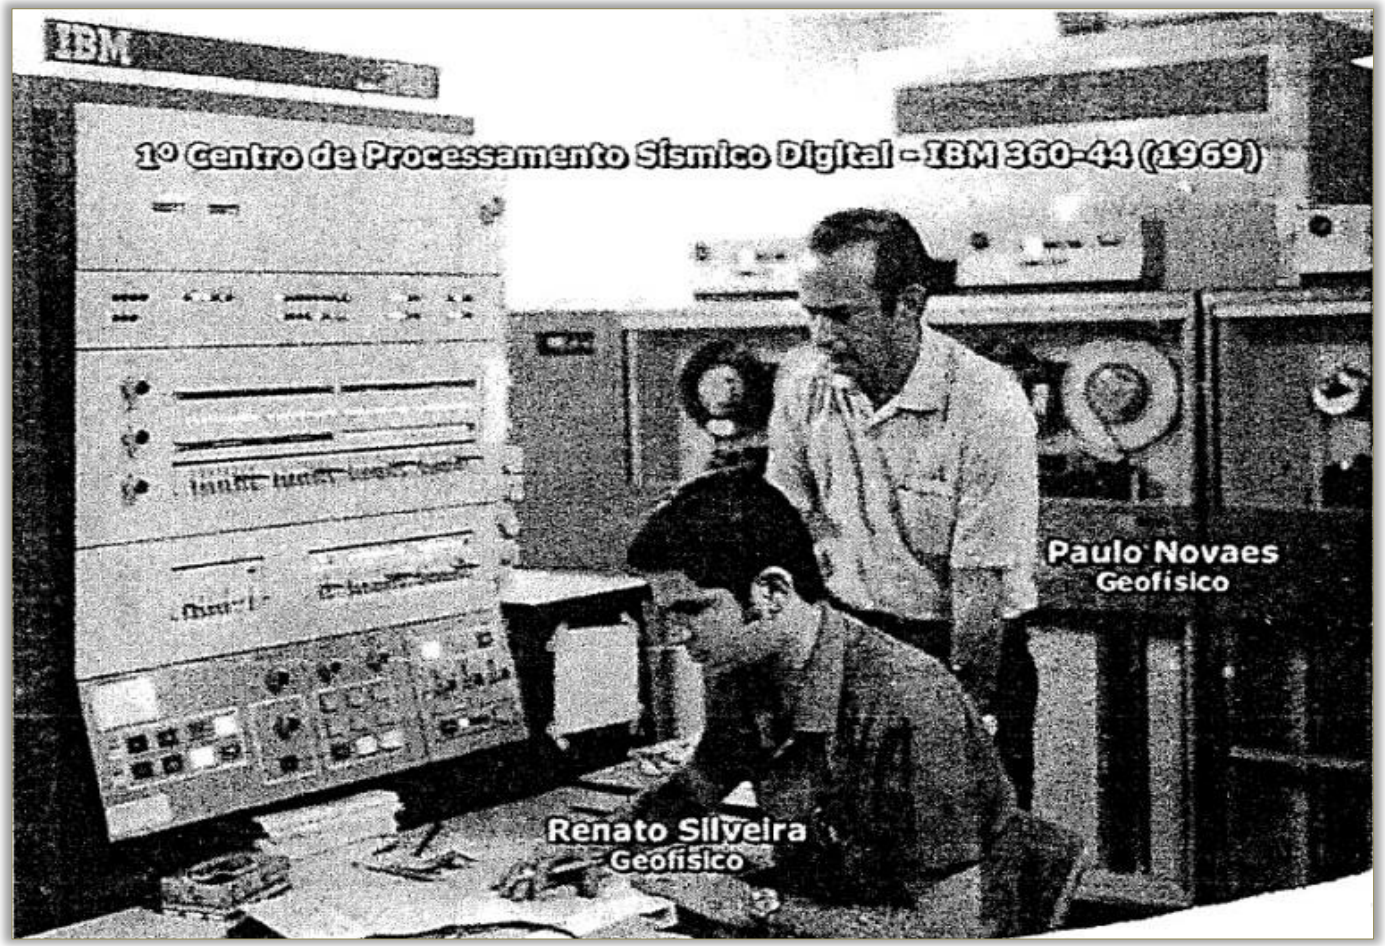

Petrobras HPC

- From 1969 to mid 1990s: IBM mainframes
- From the early 1990s: RISC machines (IBM/SGI/Sun)
- From 1997: Beowulf (x86 + Linux)

1999: Second Petrobras Beowulf
72 CPUs

2004: bwr1
1300 CPUs

2005: bwr2
4656 CPU cores

2007/2008: PS3/Cell vs Nvidia/GPU evaluation

X

- Pre-salt discovery pushed new architectures
- 4 PS3 cluster
- Port of a seismic imaging code
  - *Kirchhoff* migration

Panetta, J., Filho, P.R., Filho, C.A., Motta, F.M., Pinheiro, S.S., Junior, I.P., Rosa, A.L., Monnerat, L.R., Carneiro, L.T., & Albrecht, C.H. (2007). Computational Characteristics of Production Seismic Migration and its Performance on Novel Processor Architectures. *19th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD'07)*, 11-18.

grifo01 (2008): First Petrobras GPU cluster

- One rack in 2008, expanded to 3 racks in 2009
  - GPUs were not available in AWS and Azure till 2016
- For Kirchhoff Seismic Imaging, it was equivalent to a 4000 node CPU cluster
- For Kirchhoff Seismic Imaging: 10x gains in terms of:
  - performance/cost
  - performance/Watts
  - performance/m2
- Very good FP32, but poor FP64 performance
  - It was not ranked in the TOP500

grifo04 (2012)

- Biggest supercomputer in the O&G industry in TOP500
- Biggest supercomputer in Latin America in TOP500
- 1,088 M2050 GPUs in 544 nodes with InfiniBand DDR

Crisis....

- After having the biggest TOP500 energy supercomputer in 2012, in 2018 our HPC infrastructure was almost completely obsolete
- In 2018 we built a plan to revamp our HPC infrastrucure in waves, which led to some fragmentation
- Along the way this plan had to be modified, due to COVID-19 and other issues

Agenda

- Why we use HPC at Petrobras
- How HPC evolved at Petrobras
- Our current HPC infrastructure
- Cloud
- Conclusions

CENPES

Petrobras Research Center

- Largest applied research center in Latin America
- Located in Rio de Janeiro, Brazil, part of UFRJ campus
- Founded in 1963

PETROBRAS

HPC Datacenters

- Proprietary DC
  - +50% of its capacity is dedicated to HPC
    - Seismic
    - Reservoir
    - Machine Learning
    - Research
  - ~10 MVA capacity
  - ~2,000 square meters

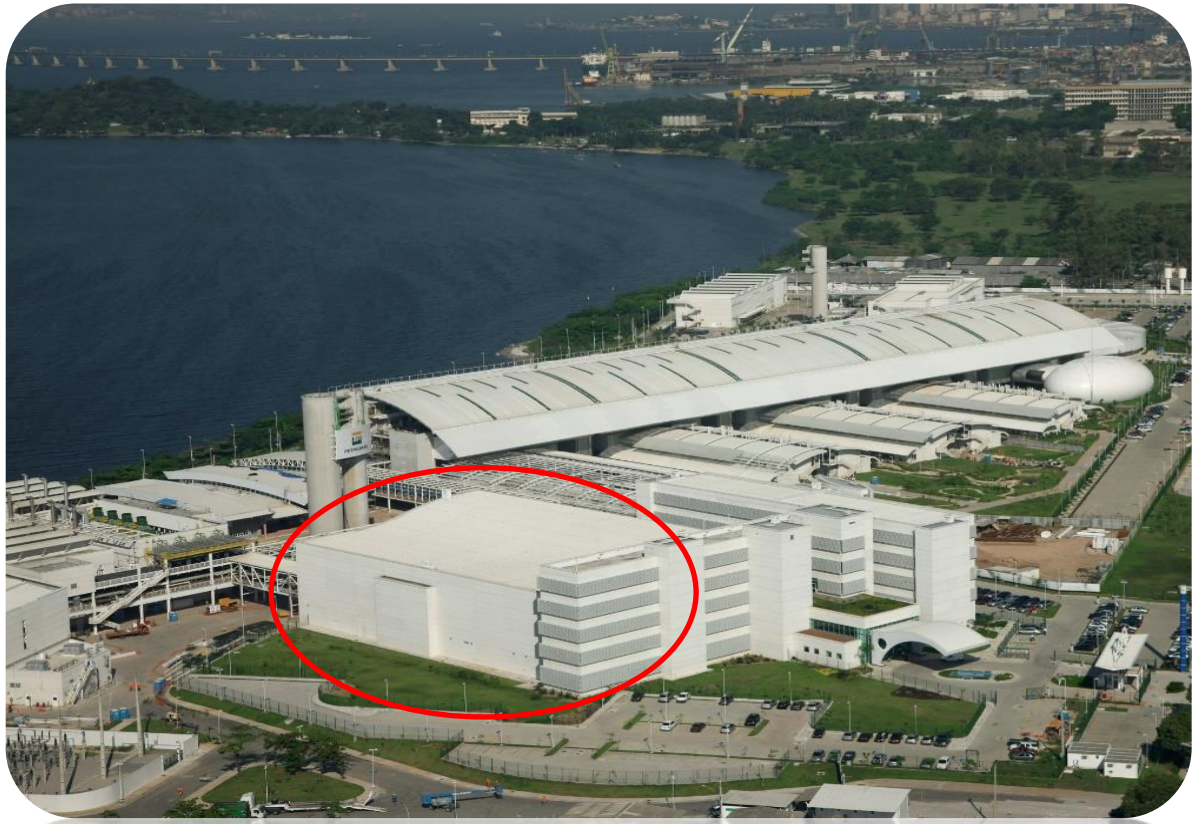

- Collocation DC
  - 95% of its capacity is dedicated to HPC
    - Mainly utilized for seismic

- Both are located in Rio de Janeiro, Brazil

Petrobras Research Center Clusters

- Atena (2020)
  - 192 GPU nodes with V100 32GB
  - 256 CPU nodes
  - InfiniBand EDR
  - 2.3 Pflops
  - Multi-purpose research
  - Singularity available
  - SLURM + small Windows HPC Pack
- Gaia (being installed[1])
  - 704 NVidia A100
  - 7.7 Pflops
  - Geophysics ML research

1 https://www.agenciapetrobras.com.br/pt/inovacao/petrobras-monta-supercomputador-para-desenvolver-tecnologias-18-01-2023/

Petrobras Reservoir Simulation clusters

Guaricema

Libra

Búzios

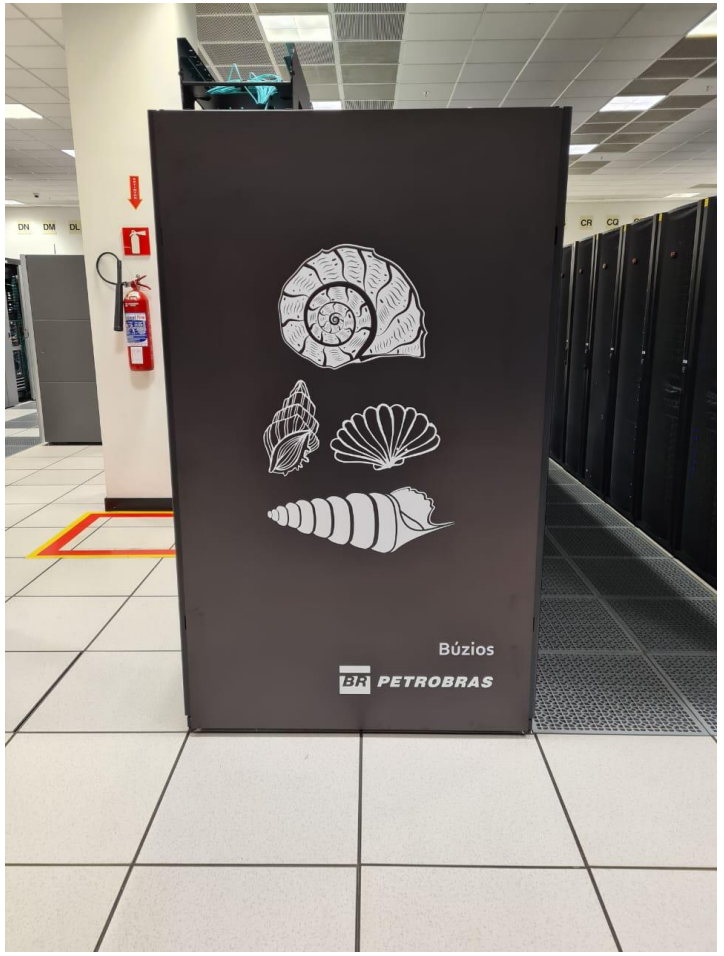

- Total of 1200 nodes, each with 40+ cores and 384GB RAM
  - Being installed: + 448 nodes with 64 cores (Albacora)
- All manage by a single SLURM scheduler
- We don't use GPUs for Reservoir Simulation (yet)

Main Seismic clusters

2019: 5 PFlops

2020: 9 PFlops

Petrobras held the Latin America TOP1 supercomputer in both the TOP500 and Green500 lists for four consecutive years.

2021: 14 PFlops

2022: 42 PFlops

Pégaso

252 nodes with Infiniband HDR

2,016 NVidia A100/80GB SXM

#33 in TOP500 - Biggest in Latin America

42 PetaFLOPS Rpeak

20 PB in Lustre filesystems

Petrobras HPC capacity evolution

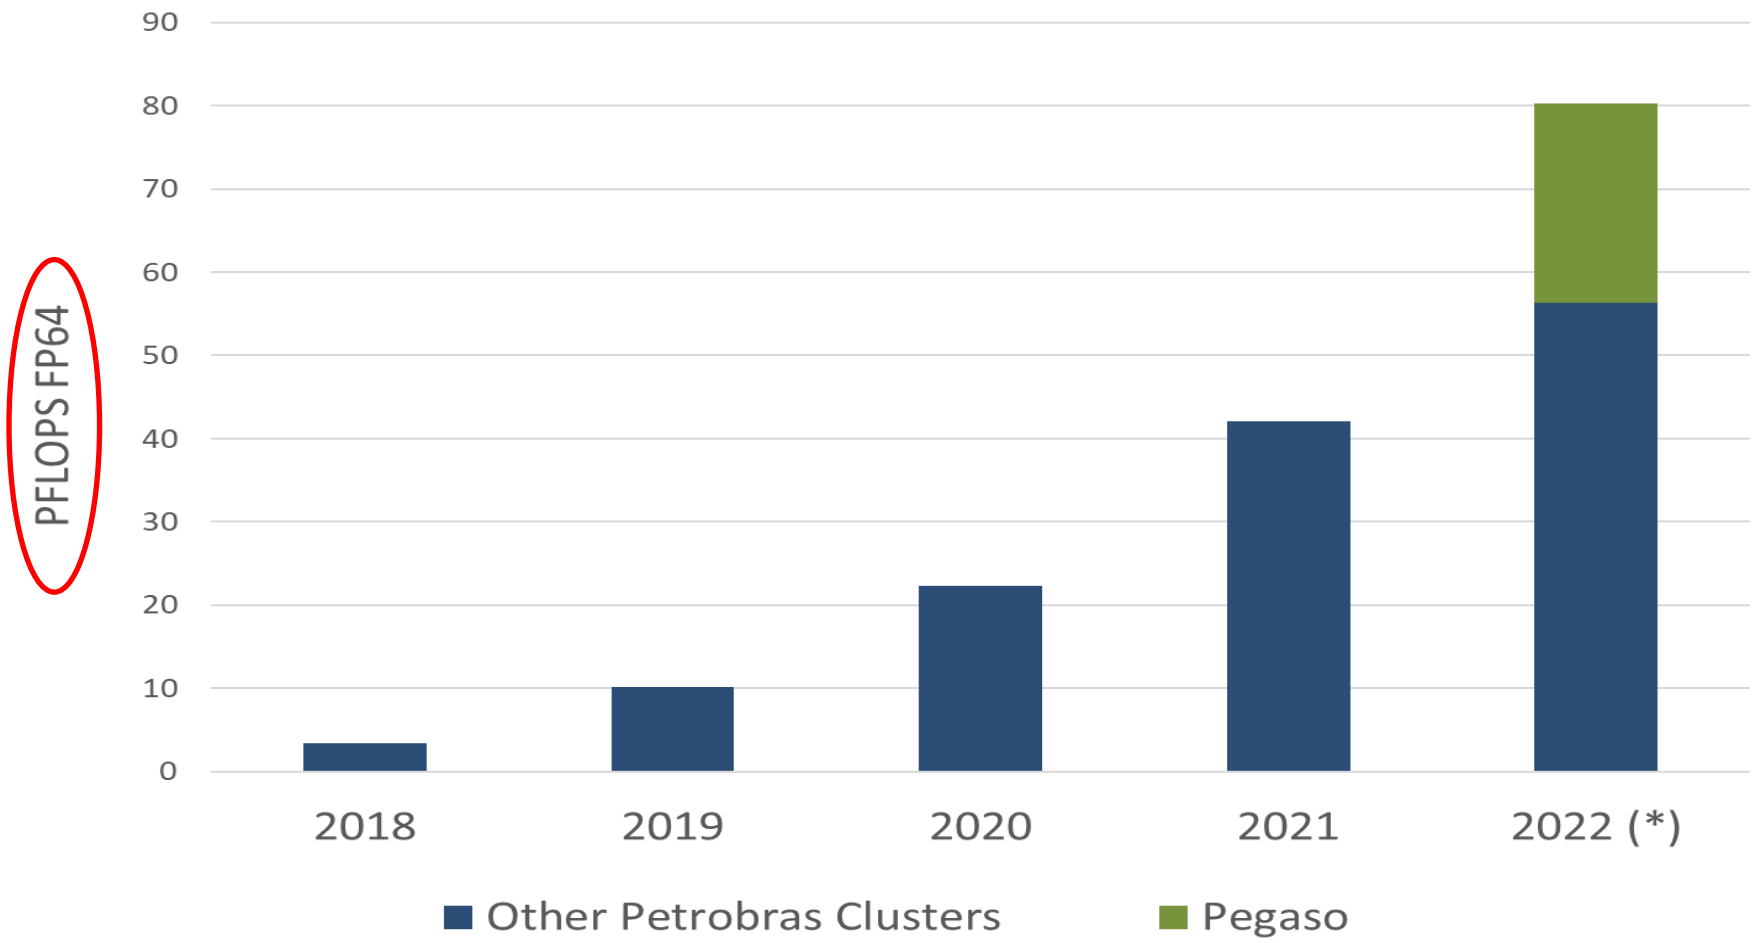

Sum of the capacity of **all 10+ Petrobras clusters**, including those not listed in the TOP500 in **FLOPS peak FP64**

O&G in the TOP500

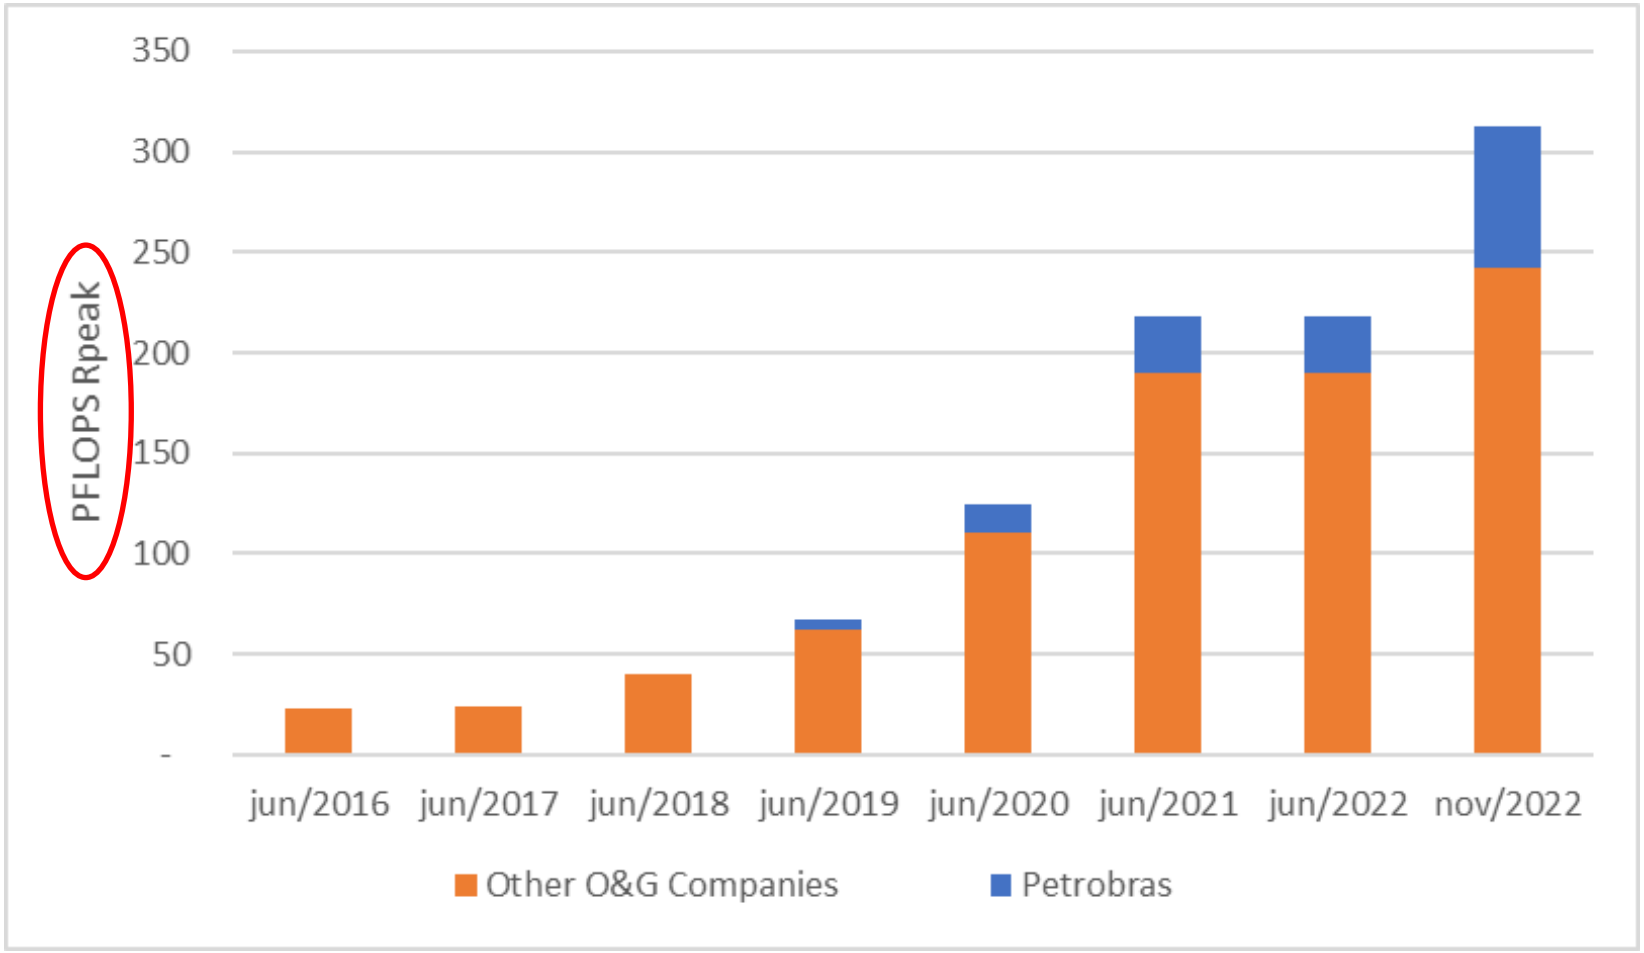

O&G vs BSC in the TOP500

PETROBRAS

MareNostrum 4
13,9 PFlops
jun/2017

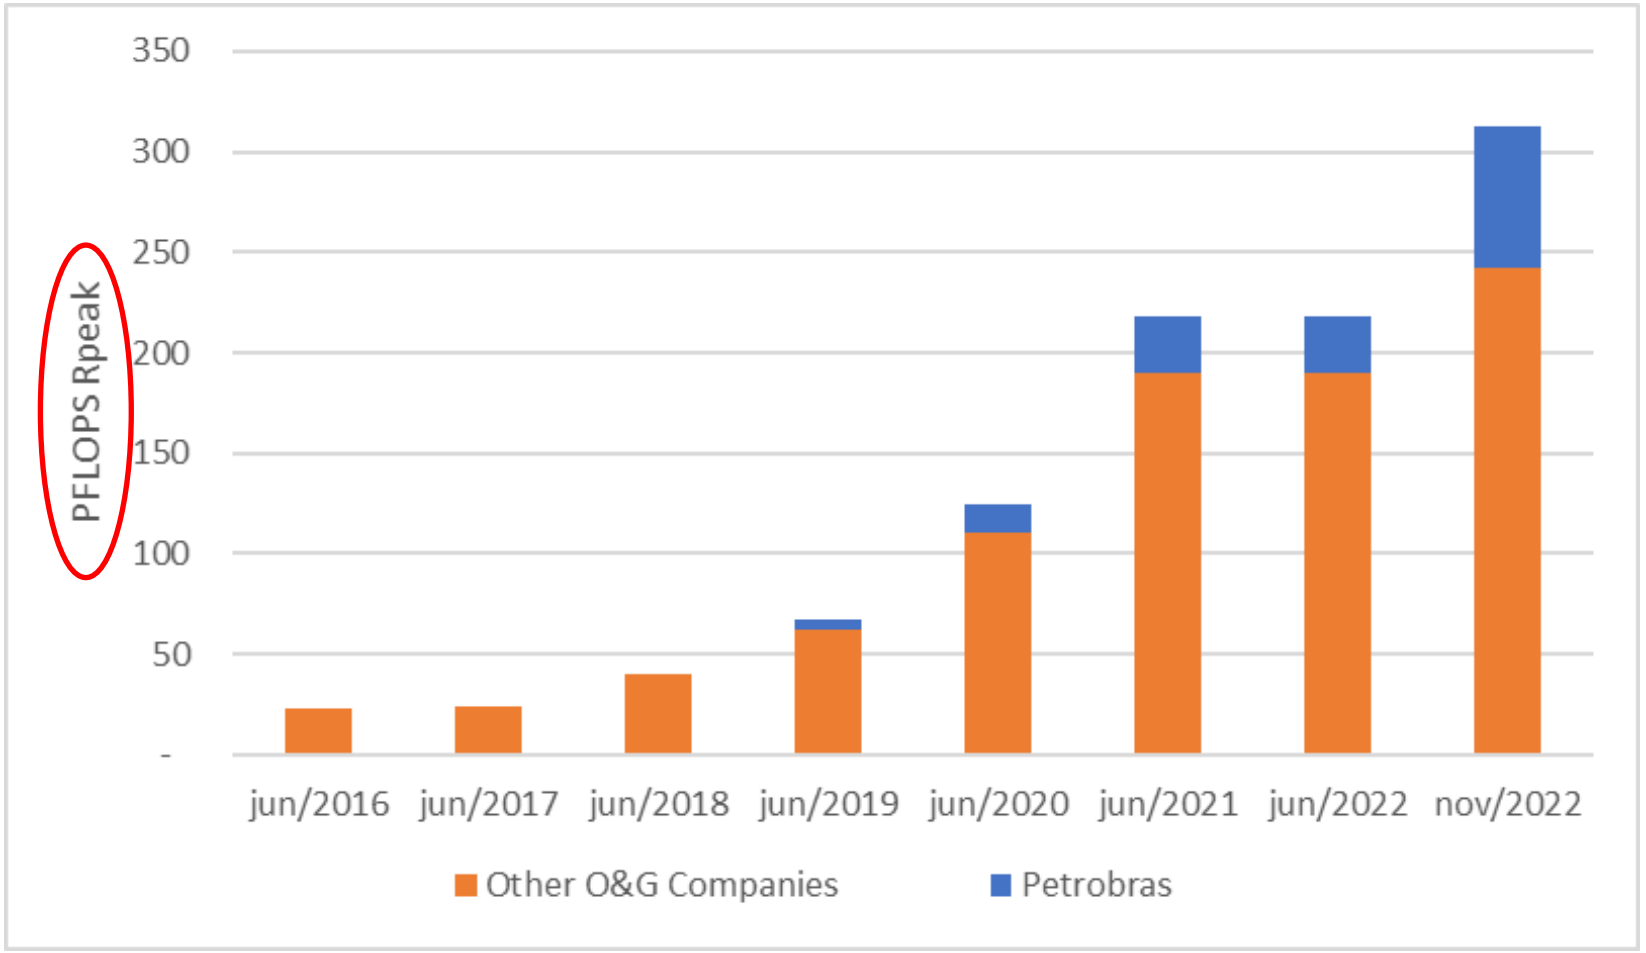

MareNostrum 5
314 PFlops
dec/2023

Academic Partnerships

- Investment in academic partners (2019)
  - 5,1 PFlops in SDumont
    - + 350 users from 50 projects from different institutions all over Brazil
  - 1,6 PFlops in OGBON
    - Geophysics Research

Agenda

- Why we use HPC at Petrobras
- How HPC evolved at Petrobras
- Our current HPC infrastructure
- Our future trends
- Conclusions

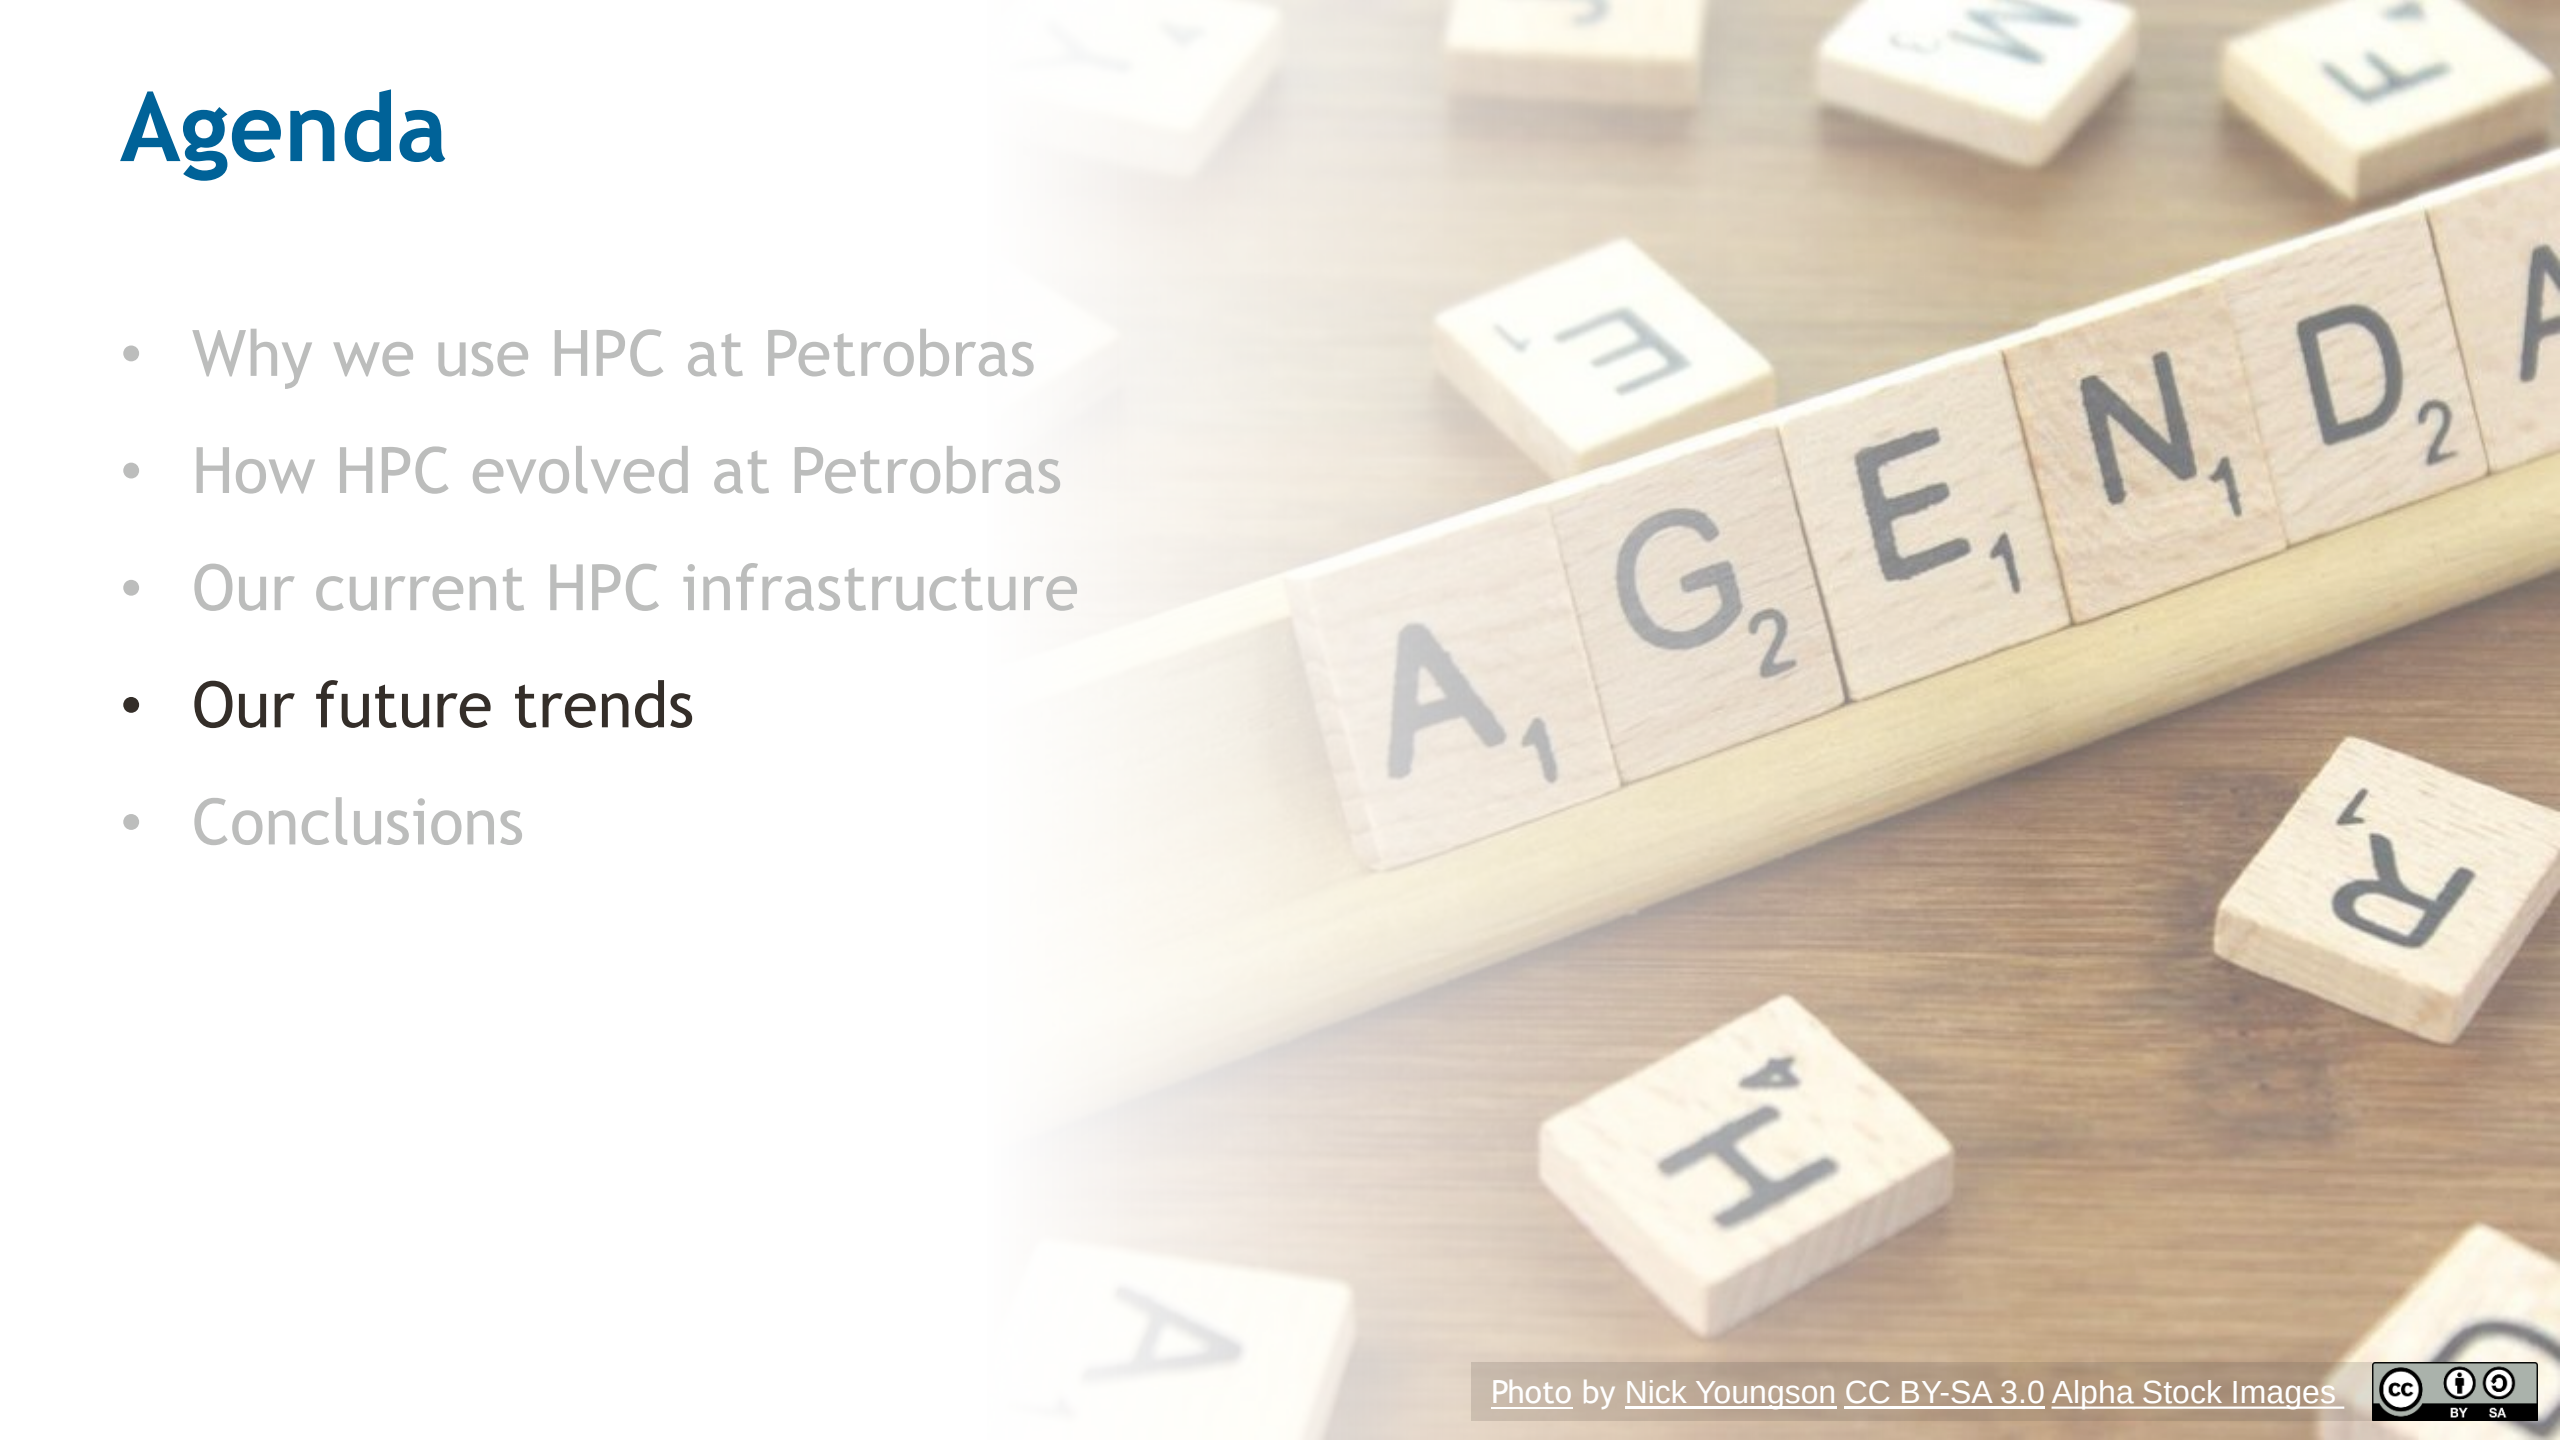

Our Future Trends: Cloud

- Cloud brings new opportunities for HPC:
  - But, for HPC it is quite more expensive than on-prem clusters if used without care
  - Has some challenges for seismic
  - More suitable for Reservoir and Machine Learning
  - Fast access to ML accelerators such Google TPUs, AWS Trainium, etc.

- Part of our HPC ecosystem since 2021
  - First steps
    - Mix of on-prem and Cloud for ML
      - SLURM on-demand clusters
    - Offload peak reservoir usage to the Cloud
      - 600 jobs in parallel

- Next steps:
  - Data cache
  - Spot usage
  - Multi-cloud

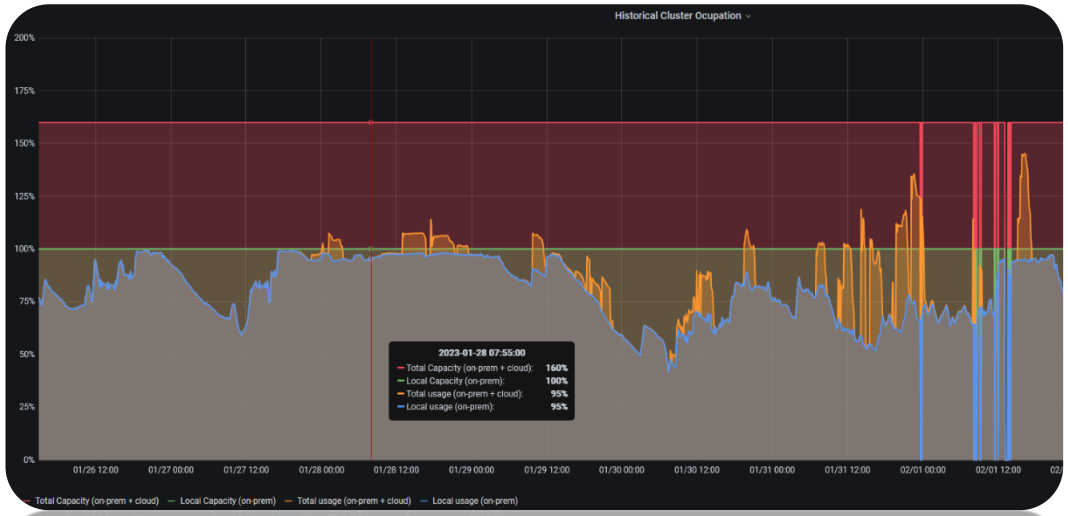

Our Future Trends: Quantum

- Currently under investigation to be quantum-ready
  - Some of our workloads may be appropriate
  - Partnerships with academic institutions
  - Prospecting algorithms using AWS Bracket
  - Many technical challenges still to be addressed

```
In [4]: #-------------------------------------------------------------------
        # FUNCAO QUE MONTA O CIRCUITO PARA CALCULAR A QFT DE 'n' QUBITS:
        #-------------------------------------------------------------------
        def QFT(n):
            F_NAME = '\'QFT()\'';
            from qiskit import QuantumRegister, ClassicalRegister, QuantumCircuit;
        #-----------------------------
            Q  = QuantumRegister(  n, name = 'm');
            C  = ClassicalRegister(n, name = 'c');
            QC = QuantumCircuit(Q, C, name = "QFT");
        #-----------------------------
            for i in range(0, n):
                m = n - 1 - i;
                QC.h(Q[m]);
                for j in range(m, 0, -1):
                    k = 2 + m - j;
                    Rk(QC, Q[j-1], Q[m], k, 1.0);
                QC.barrier();
        #-----------------------------
        # PORTA SWAP:
        #-----------------------------
            j = 0;
            while ( int(n/2) > j ):
                QC.swap( Q[j], Q[n-1 - j] );
                j = j + 1;
        #-----------------------------
            QC.barrier();
            QC.measure(Q, C);
            print(' ---------- FUNCAO %s FUNCIONA---------- '% F_NAME);
        #-----------------------------
            return QC;
        #-------------------------------------------------------------------
        # FUNCAO QUE DEFINE A PORTA Rk USANDO A PORTA cu1:
        #-------------------------------------------------------------------
        def Rk(QC, Q_CONTROL, Q_TARGET, k, a):
```

Quantum algorithm on AWS. Courtesy: Eldues Martins

ÚLTIMAS NOTÍCIAS

Pesquisadores do CBPF e UNICAMP desenvolvem um chip quântico supercondutor

Publicado: Terça, 22 de Fevereiro de 2022, 14h49 Última atualização em Terça, 22 de Fevereiro de 2022, 14h49 Acessos: 699

Tweetar Curtir 52 m

Um chip quântico supercondutor foi desenvolvido por pesquisadores do Centro Brasileiro de Pesquisas Físicas (CBPF) e da Universidade Estadual de Campinas (UNICAMP). O CBPF conta com importante apoio da Petrobras e da Fundação de Amparo à Pesquisa do Estado do Rio de Janeiro (FAPERJ), por meio do Programa Cientista do Nosso Estado. Já o grupo do Instituto de Física Gleb Wataghin (IFGW) da Unicamp recebe apoio da Fundação de Amparo à Pesquisa do Estado de São Paulo (FAPESP).

Os computadores quânticos estão no topo da agenda científico-tecnológica dos países desenvolvidos que, juntos, investiram no ano de 2020 mais de US$ 20 bilhões no desenvolvimento de tecnologias quânticas. Por este motivo, ao longo da última década, importantes progressos foram alcançados no desenvolvimento de bits quânticos (q-bits) supercondutores, como uma arquitetura escalável para o processamento de informação quântica. Devido a estes avanços, empresas como as norte-americanas IBM, Google e Rigetti, bem como a canadense D-Wave, apostam na tecnologia de circuitos supercondutores para a fabricação de chips de computadores quânticos.

O computador quântico

O computador, seja ele clássico ou quântico, é um objeto físico, portanto sujeito às leis da Física. Para entender

Source: CBPF Website

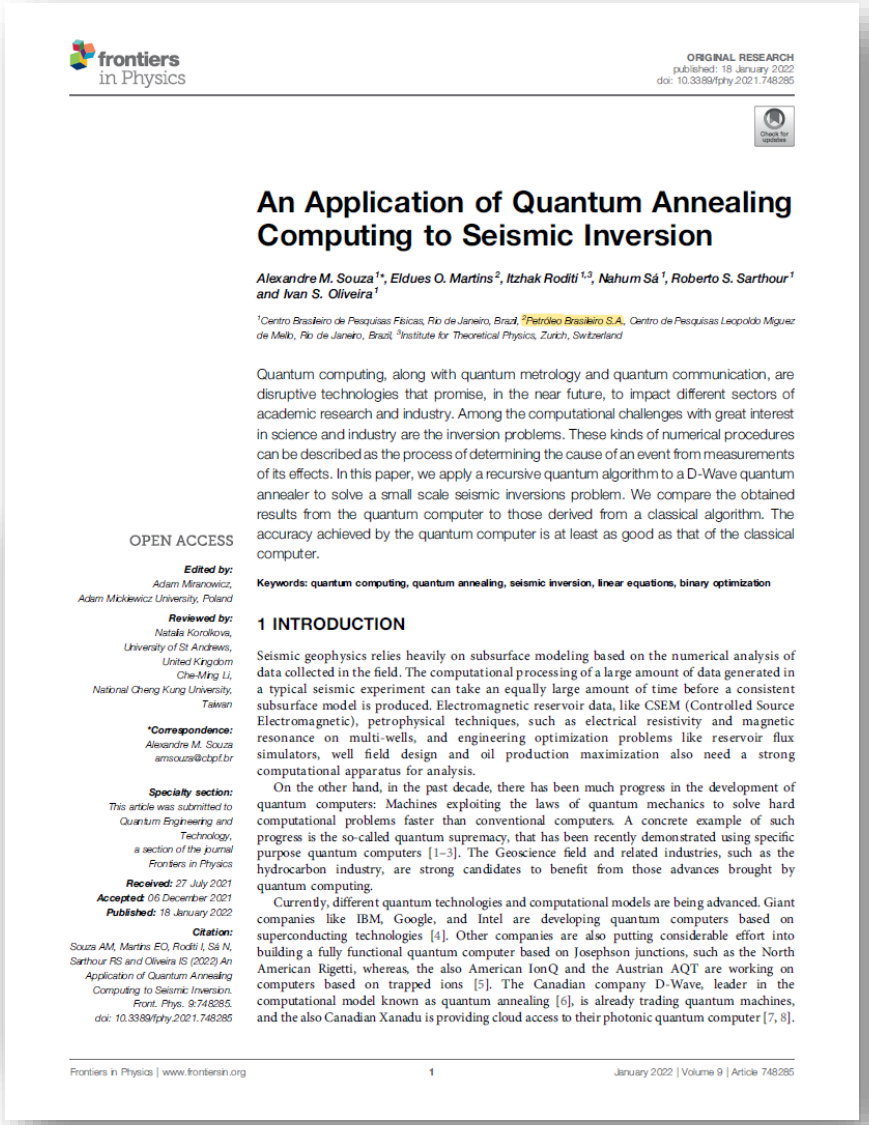

frontiers in Physics

ORIGINAL RESEARCH
published: 18 January 2022
doi: 10.3389/fphy.2021.748285

An Application of Quantum Annealing Computing to Seismic Inversion

Alexandre M. Souza[1], Eldues O. Martins[2], Itzhak Roditi[1,3], Nahum Sá[1], Roberto S. Sarthour[1] and Ivan S. Oliveira[1]*

[1]Centro Brasileiro de Pesquisas Físicas, Rio de Janeiro, Brazil, [2]Petróleo Brasileiro S.A., Centro de Pesquisas Leopoldo Miguez de Mello, Rio de Janeiro, Brazil, [3]Institute for Theoretical Physics, Zurich, Switzerland

Quantum computing, along with quantum metrology and quantum communication, are disruptive technologies that promise, in the near future, to impact different sectors of academic research and industry. Among the computational challenges with great interest in science and industry are the inversion problems. These kinds of numerical procedures can be described as the process of determining the cause of an event from measurements of its effects. In this paper, we apply a recursive quantum algorithm to a D-Wave quantum annealer to solve a small scale seismic inversions problem. We compare the obtained results from the quantum computer to those derived from a classical algorithm. The accuracy achieved by the quantum computer is at least as good as that of the classical computer.

OPEN ACCESS

Edited by:
Adam Miranowicz, Adam Mickiewicz University, Poland

Reviewed by:
Natalia Korolkova, University of St Andrews, United Kingdom
Che-Ming Li, National Cheng Kung University, Taiwan

***Correspondence:**
Alexandre M. Souza
amsouza@cbpf.br

Specialty section:
This article was submitted to Quantum Engineering and Technology, a section of the journal Frontiers in Physics

Received: 27 July 2021
Accepted: 06 December 2021
Published: 18 January 2022

Citation:
Souza AM, Martins EO, Roditi I, Sá N, Sarthour RS and Oliveira IS (2022) An Application of Quantum Annealing Computing to Seismic Inversion. Front. Phys. 9:748285. doi: 10.3389/fphy.2021.748285

Keywords: quantum computing, quantum annealing, seismic inversion, linear equations, binary optimization

1 INTRODUCTION

Seismic geophysics relies heavily on subsurface modeling based on the numerical analysis of data collected in the field. The computational processing of a large amount of data generated in a typical seismic experiment can take an equally large amount of time before a consistent subsurface model is produced. Electromagnetic reservoir data, like CSEM (Controlled Source Electromagnetic), petrophysical techniques, such as electrical resistivity and magnetic resonance on multi-wells, and engineering optimization problems like reservoir flux simulators, well field design and oil production maximization also need a strong computational apparatus for analysis.

On the other hand, in the past decade, there has been much progress in the development of quantum computers: Machines exploiting the laws of quantum mechanics to solve hard computational problems faster than conventional computers. A concrete example of such progress is the so-called quantum supremacy, that has been recently demonstrated using specific purpose quantum computers [1–3]. The Geoscience field and related industries, such as the hydrocarbon industry, are strong candidates to benefit from those advances brought by quantum computing.

Currently, different quantum technologies and computational models are being advanced. Giant companies like IBM, Google, and Intel are developing quantum computers based on superconducting technologies [4]. Other companies are also putting considerable effort into building a fully functional quantum computer based on Josephson junctions, such as the North American Rigetti, whereas, the also American IonQ and the Austrian AQT are working on computers based on trapped ions [5]. The Canadian company D-Wave, leader in the computational model known as quantum annealing [6], is already trading quantum machines, and the also Canadian Xanadu is providing cloud access to their photonic quantum computer [7, 8].

DOI: https://doi.org/10.3389/fphy.2021.748285

Our Future Trends: HPCaaS
HPC as a Service for Reservoir Simulation

- Many recent changes in Reservoir Engineer ecosystem
  - GPU and MPI simulators becoming available commercially and in-house
    - Not all models fit them
  - Cloud has instance offers that could be suitable to different models
  - It presents challenges to the engineer
    - Decide the best configuration to choose
    - Best for their model vs Best globally

- Research project
  - ➢ Focus the reservoir engineers' jobs on SLA
  - ➢ Being agnostic to the underlying HPC infrastructure
    - ➢ On-premises vs cloud
    - ➢ Instance type
    - ➢ GPU vs CPU

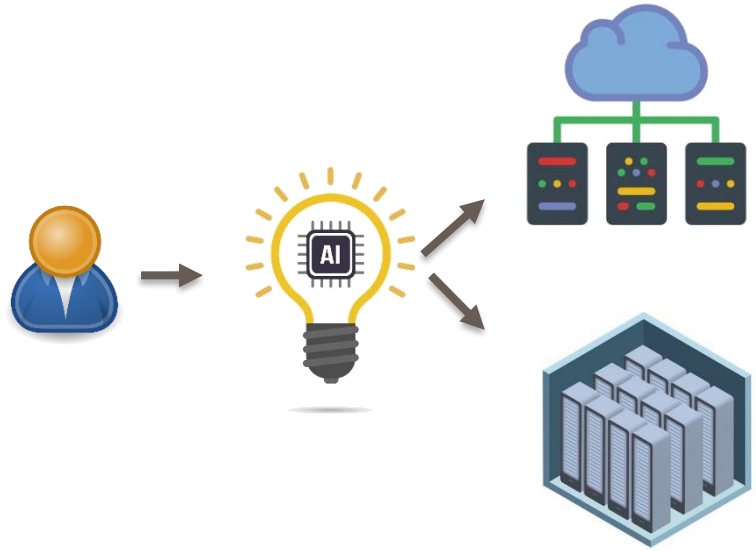

Conclusions

- HPC is very important for Petrobras business
  - Specially for seismic processing and reservoir simulation
- We started with “HPC” back in 1969
  - Of course it evolved a lot since then
  - First x86 Linux cluster in 1997
  - Since 2008 we’ve been using GPUs in production
- Since 2018 we expanded our HPC capacity by more than 20x
  - Internally: 10+ clusters
  - Partners: 2 academic supercomputers
  - We will keep growing
- Investing in the Cloud and in ways to make better use of it

THANK YOU!!!

felipeportella@petrobras.com.br

linkedin.com/in/felipeportella

Pégaso Video