

**the
human
energy
company**[®]



HPC in Chevron A journey to the cloud

Alex Loddoch
Technology Strategist
Chevron Fellow

This document is intended only for use by Chevron for presentation at the 89th HPC User Forum and inclusion by the Hyperion Research on a conference website that is available to attendees of the HPC User Forum. No portion of this document may be copied, displayed, distributed, reproduced, published, sold, licensed, downloaded, or used to create a derivative work, unless the use has been specifically authorized by Chevron in writing.

High Performance Computing

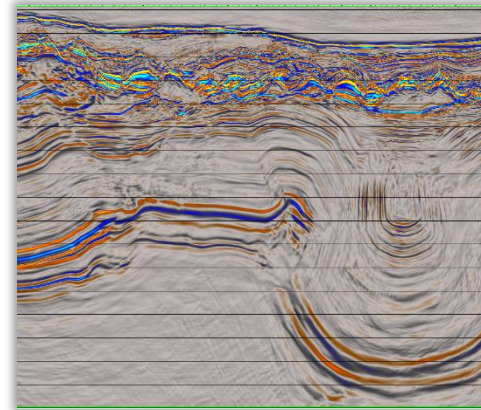
why do we care?

exploration success depends on subsurface insight



seismic survey

petabytes of raw data



seismic imaging

high-impact decision making



deep water well

reducing uncertainty

**Single update to subsurface image
100,000,000,000,000,000,000 operations = 100 Exaflop**

Laptop
years

High-end PC
months

Chevron HPC
hours

LLNL El Capitan (#1)
minutes

High Performance Computing in Chevron

where do we use it?

Seismic processing, seismic imaging

Process data, reconstruct images
Range of methods and algorithms

Computational geology

Study formation of reservoirs at grain scale
Multiphysics (CFD + sedimentation/erosion)

Reservoir simulation

Simulate/optimize processes within a reservoir
Multiphysics (multiphase flow + chemistry)

Rock mechanics, structural mechanics




Study fracturing of rocks under pressure
Finite element methods

Data science, machine learning, AI

growing capability growing the corporation
Random forest, CNN, ML/AI, GPT, RAG, ...


Downstream

Optimizing refinery operations (linear programming)
Developing catalytic materials (molecular modeling)



long history of HPC and technology innovation

original users of “supercomputers”

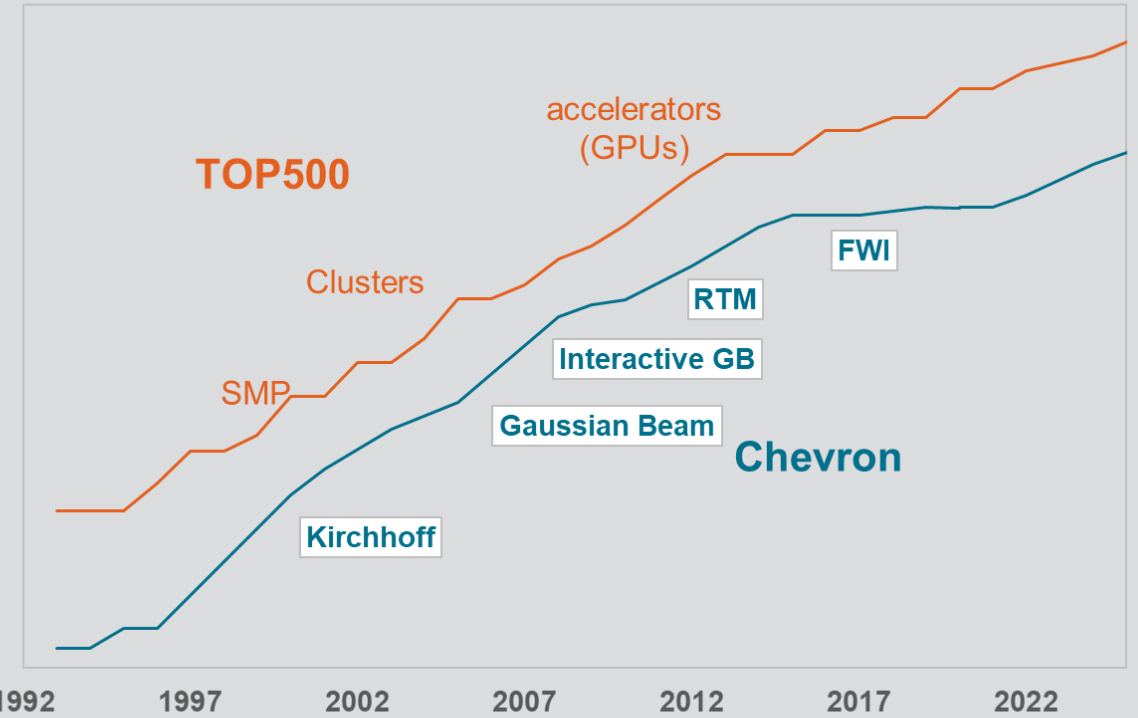


pioneering the use of GPUs



continued innovation driven by business needs

evolution of compute performance



source: TOP500.org and Chevron-internal data



HPC in the cloud

a 20,000-foot view from up there



Chevron cloud HPC - a timeline

HPC cloud investigations since ~2014

Chevron opens an enterprise-scale datacenter in 2016

Chevron & Microsoft Azure enter a strategic partnership in 2017

lift & shift of production HPC environment to Azure begins in 2021

majority of subsurface workflows migrated to Azure in 2022

Chevron has ongoing efforts to cloud-naturalize workflows

lookback to 2018

what did we think cloud had to offer for HPC?

technology

scale

service structure

decision – deference

hype

what cloud has to offer for HPC

“wrong” hardware

specific HPC technologies missing, locally (IB, GPUs, ...) or generally (PFS)

additional complexity and lack of control

opaque/cloudy layer for IaaS, orchestration, data transfer/management

lack of HPC expertise

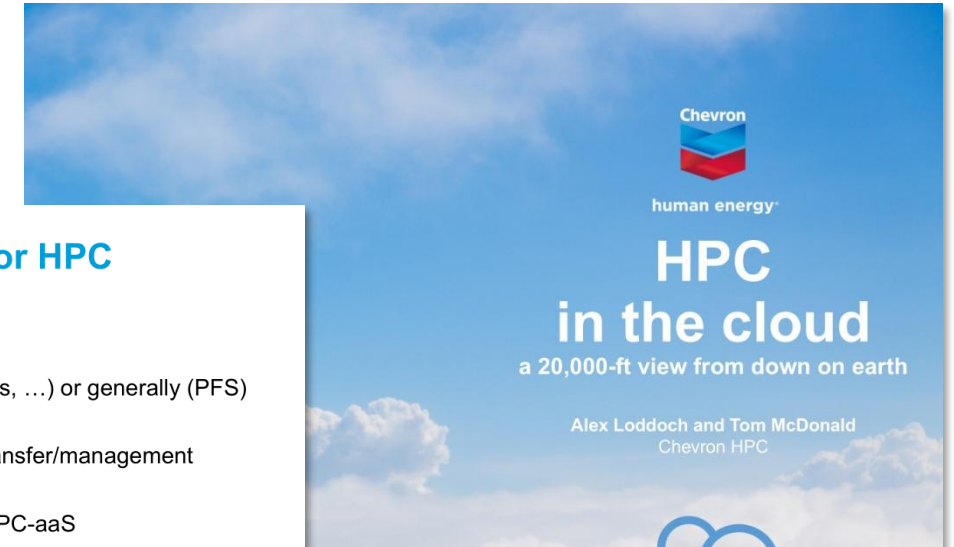
very little HPC experience with cloud vendors, no HPC-aaS

hype

“cloud resources are infinite and free”, “microservice-everything”

There is a significant difference
(but often a poor decision)
public cloud vs. private

© 2018 Chevron All Rights Reserved



what cloud has to offer for HPC

technology

various hardware and software technologies, some exclusive to cloud

scale

access to very large resource pools (“capacity”) (compute and/or storage)

service structure

fine-grained, consumption-based accounting and billing

hype

“cloud resources are infinite and free”, “microservice-everything”

generic cloud-typical technologies
(Spark, Containers, Tensorflow, ...)

vs.

proprietary, cloud-exclusive technologies
(AWS Lambda, Azure Functions, Google ImageNet...)

© 2018 Chevron All Rights Reserved



HPC in the cloud – current state

bulk compute makes up large fraction of the workload

cloud-providers are **HPC-aware**

Lift-and-Shift of HPC to cloud is **possible** and **useful**

work to **cloud-naturalize HPC** workloads is needed and ongoing

accounting and **user behavior** can be challenges

HPC in the cloud – the path forward

questions and concerns

(how) do we deal with data locality?

data volumes are large and growing, IO matters for some applications

how does cloud accommodate (HPC) technology evolution?

water cooling, AI/ML chips, esoteric hardware, tightly coupled workloads, ...

bulk compute is not capability compute is not frontier compute!

we need all three of them, along with tools to use them effectively

HPC in the cloud – the path forward *what do we need?*

collaboration on capabilities and solutions

the right products combined in the right way (incl. research)

HPC-grade cloud-native approaches and abstraction layers

we need "bulkMPI" and abstraction layers

honest conversation about goals (and billing)

what are we really after in the cloud?

any questions?


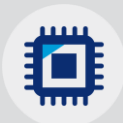

loddoch@chevron.com

#hpcmeansbusiness

appendix



HPC compute types – a differentiated approach

	description	value drivers	workload characteristics	location & partnerships
 <p>bulk compute</p>	<p>common, large capacity compute workloads with HPC requirements, consumed as services</p>	<p>efficiency</p> <p>table – stakes</p>	<p>fault-tolerant</p> <p>minimal intervention</p> <p>service-based (SaaS, PaaS)</p>	<p>cloud (location irrelevant)</p>
 <p>capability compute</p>	<p>specific HPC workloads, run in a controlled environment</p>	<p>differentiation</p> <p>control over technologies, configuration, implementation and access</p>	<p>complex architecture</p> <p>performance sensitive</p>	<p>location relevant (some IaaS)</p>
 <p>frontier compute</p>	<p>innovation and computing at performance levels or with technologies not commercially available</p>	<p>innovation</p> <p>access to capabilities at leading, pre-commercial edge of HPC</p>	<p>“Exascale”</p> <p>“can’t be done today”</p>	<p>Leadership Computing Facilities</p> <p>National Labs</p>