How Best to Leverage Cloud for (Big) HPC Sites

CUG 2025

Bill Nitzberg, Chief Scientist
Ian Littlewood, VP, HPCWorks Tools

ALTAIR

Safe Harbor Statement

This presentation and the accompanying oral commentary may contain "forward-looking" statements that are based on our beliefs and assumptions and on information available to us as of the date of this presentation. All statements other than statements of historical facts contained in this presentation, including statements regarding our future results of operations and financial position, customer lifetime value, strategy and plans, market size and opportunity, competitive position, industry environment, potential growth opportunities and our expectations for future operations, are forward-looking statements. The words "believe," "may," "might," "objective," "ongoing," "will," "estimate," "continue," "anticipate," "design," "intend," "expect," "could," "plan," "potential," "predict," "project," "seek," "should," "would" or the negative version of these words and similar expressions are intended to identify forward-looking statements. This presentation also contains non-GAAP financial measures. We have provided a reconciliation of such non-GAAP financial measures to the most directly comparable measures prepared in accordance with U.S. GAAP in the Appendix to this presentation.

We may not actually achieve the plans, intentions or expectations disclosed in our forward-looking statements. Forward-looking statements involve known and unknown risks, uncertainties, assumptions and other factors that may cause our actual results, performance, achievements or expectations to be materially different from any future results, performance, achievements or expectations expressed or implied by the forward-looking statements. Except as required by law, we assume no obligation to update these forward-looking statements publicly, or to update the reasons why actual results could differ materially from those anticipated in the forward-looking statements, even if new information becomes available in the future.

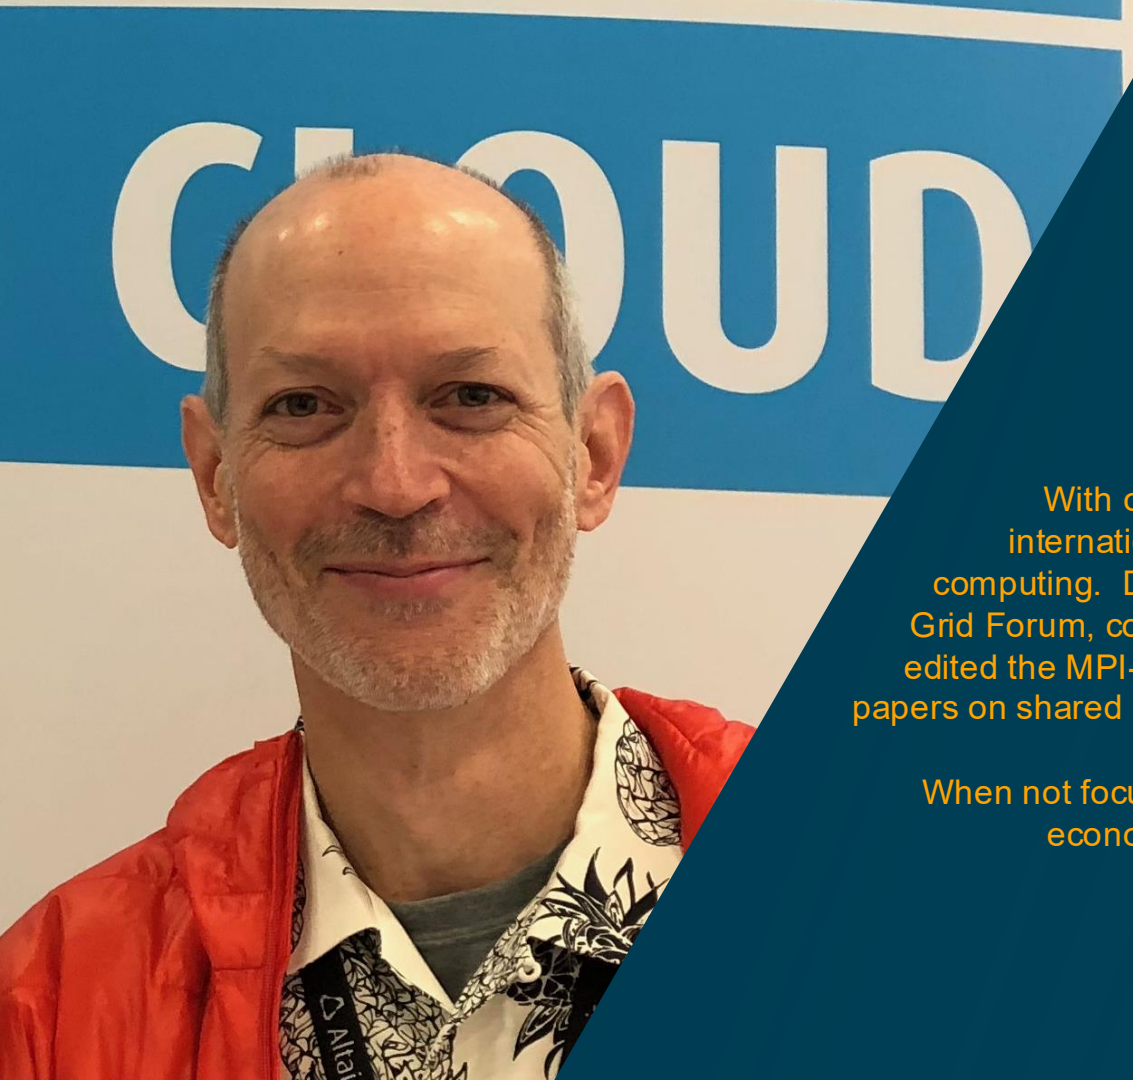

WHO IS BILL?

Bill Nitzberg, Chief Scientist, Altair

With over 30 years in computing, Dr. Nitzberg is an internationally recognized expert in high performance computing. Dr. Nitzberg served on the board of the Open Grid Forum, co-architected NASA's Information Power Grid, edited the MPI-2 I/O standard, and has published numerous papers on shared memory, parallel I/O, scheduling, and clouds.

When not focused on HPC, Bill tries to improve his running economy for his long-distance running adventures.

ALTAIR

Changing Tomorrow, Together

1985
Founded

16k+
Customers

$666M
FY2024 Revenue

Altair has been acquired by Siemens, creating the world's most complete AI-powered portfolio of industrial software for simulation, high performance computing, data science, and artificial intelligence as part of the Siemens Xcelerator platform.

ALTAIR

Data Science + Rocket Science™ =

Altair One | Cloud Innovation Gateway

ALTAIR
HYPERWORKS
Design & Simulation Platform

ALTAIR
HPCWORKS
HPC & Cloud Platform

ALTAIR
RAPIDMINER
Data Analytics & AI Platform

Altair HPCWorks Platform

Accelerate and optimize HPC, AI, and Cloud

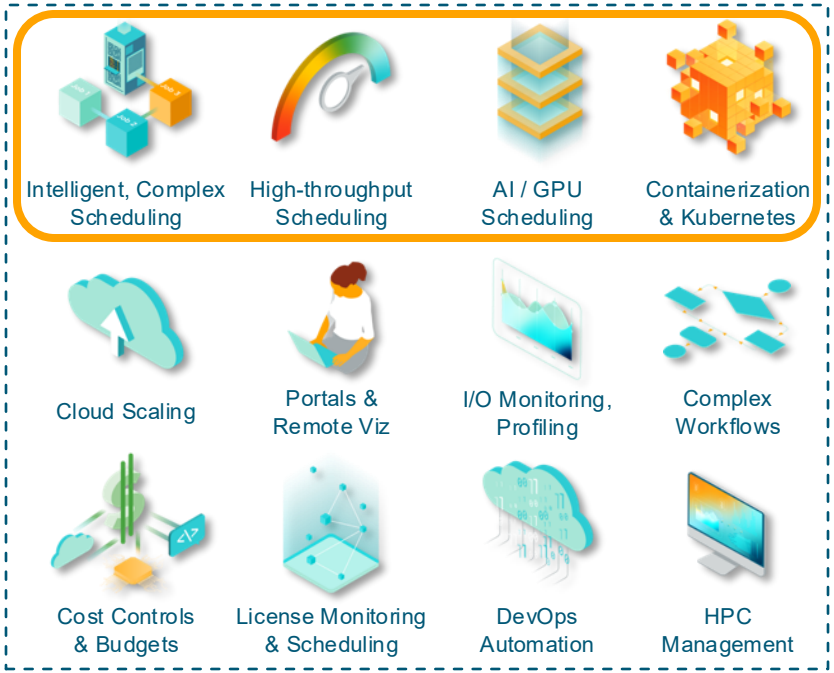

Achieving breakthrough results requires computing and analyzing more (a lot more), faster (a lot faster),

but, *it's complex*, *it's fragile*, and *it's expensive.*

HPCWorks delivers

- Easier access (a lot easier)
- Faster turnaround (a lot faster)
- Rapid Scaling™ for more compute (a lot more)
- Cost controls and visibility (a lot more control)
- Future-proof, open ecosystem (a lot more flexibility)
- Plus… optimize flows, licenses, I/O, energy, GPUs, quantum, emulators, DevOps, Cloud, … (a lot more efficiency)

ALTAIR

Quick Poll…

Which will happen in <= 5 years?

Self-driving cars (for a lot of your car rides)

Fusion energy (for a lot of our energy)

Quantum computing (for a lot of the TOP500)

AGI - Artificial General Intelligence

Cloud “clusters” on TOP500 (for a lot of the TOP500)

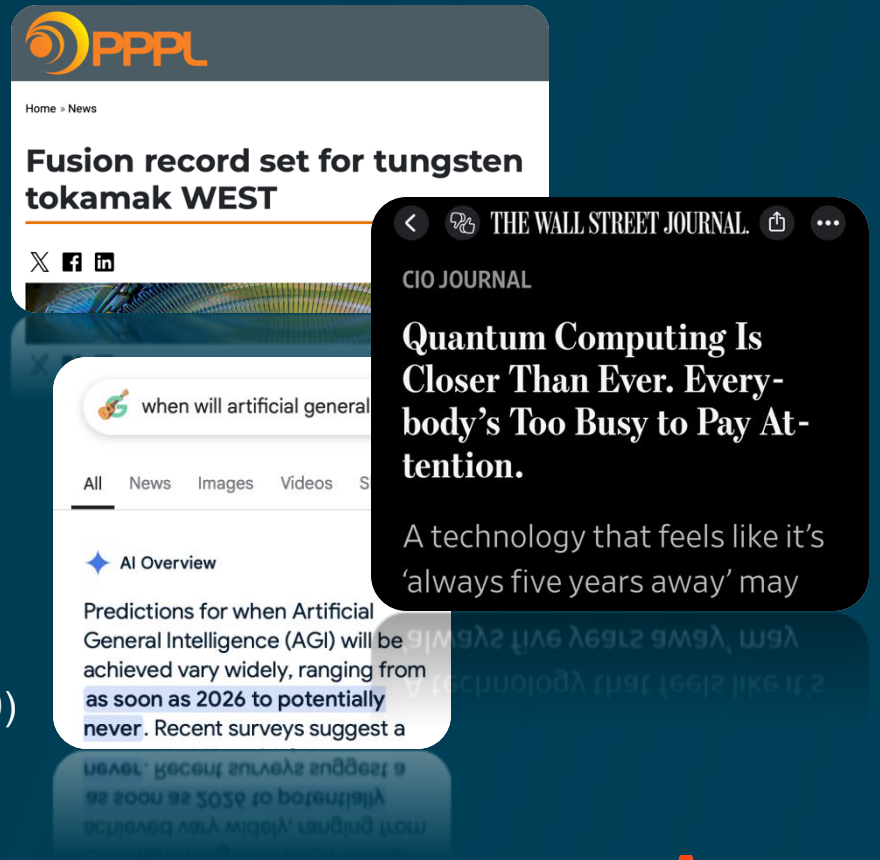

ALTAIR

WHAT IS CLOUD?

Cloud = Someone Else's Computer (Rented)

$664,000 / 50 years = $1.52/hr

$3099 / 3.5 years = $0.10/hr

7-25x more $

$280 / 1 day = $11.67 / hr

aws

$0.55-2.54 / hr

ALTAIR

WHY CLOUD?

ALTAIR

Cloud is a Panacea…

Infinite compute → Infinite Profit

HYPERION RESEARCH

"For every $1 spent on HPC, businesses see $463 in incremental revenues and $44 in incremental profit"

On-premise – constrained

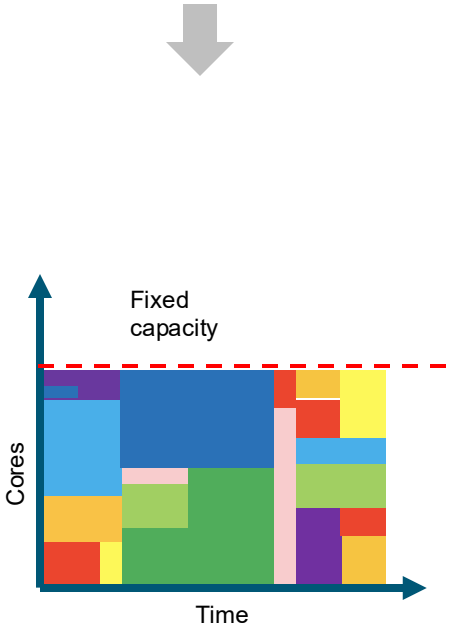

Cloud – infinite compute

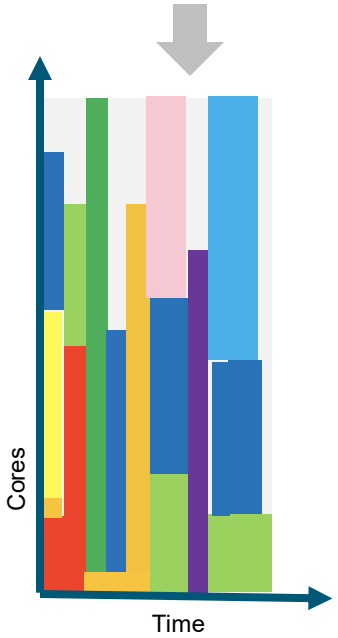

ALTAIR

Circa 2011…

"Infinite capacity – get results ASAP!"

"99.999% reliability"

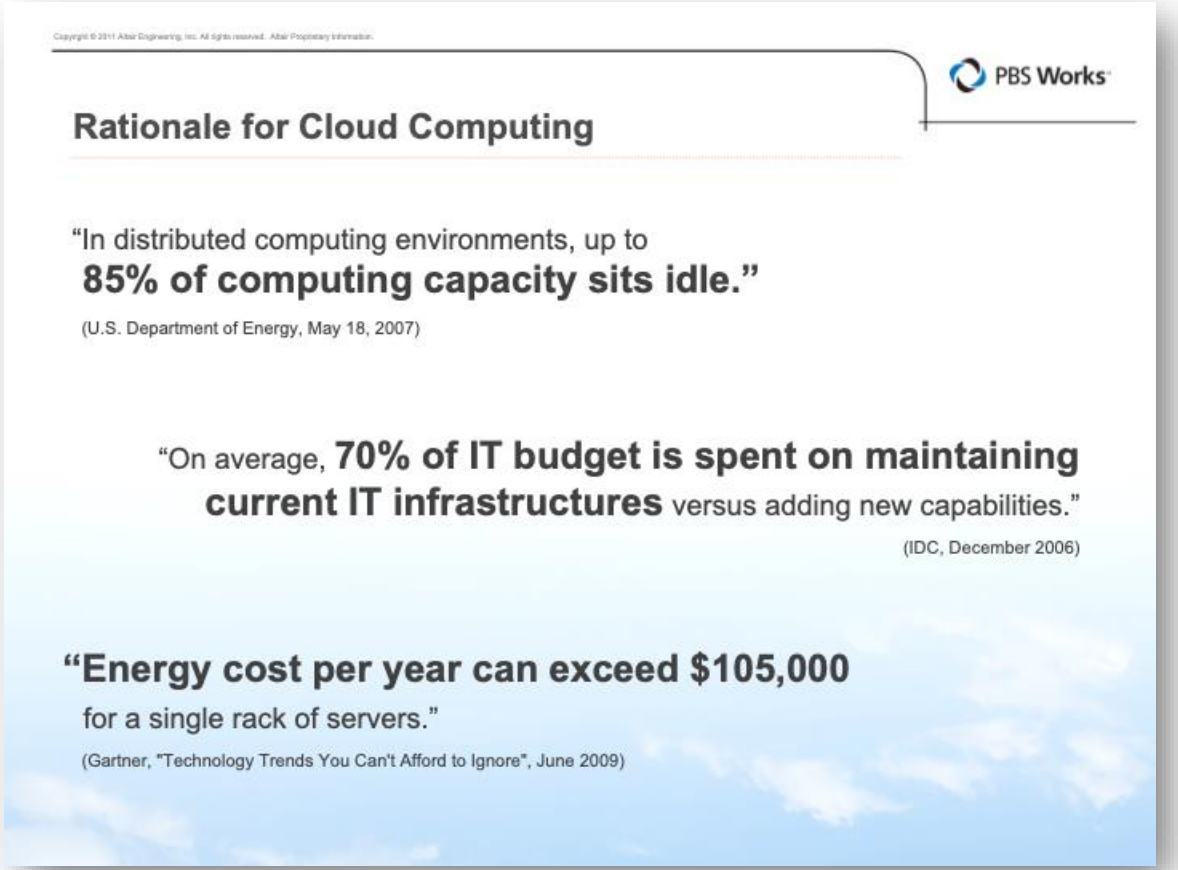

"Save $$$ (only pay for what's used)"

"Shrink you IT staff"

"Cloud has economy-of-scale"

ALTAIR

But…

"~~Infinite capacity~~ – get results ASAP!"

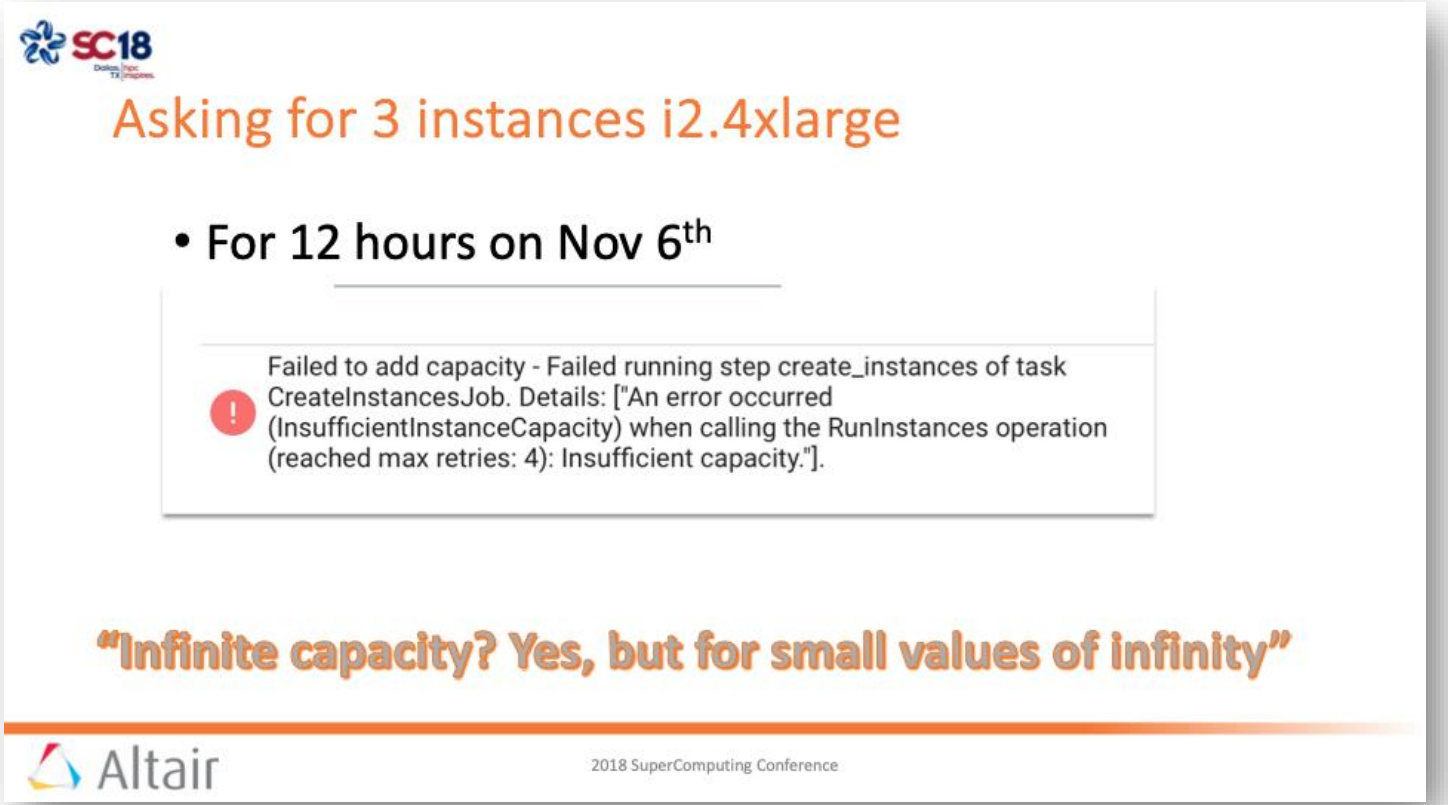

ALTAIR

Scorecard vs. Circa 2011

~~"Infinite capacity – get results ASAP!"~~

~~"99.999% reliability"~~

Data is still a big issue

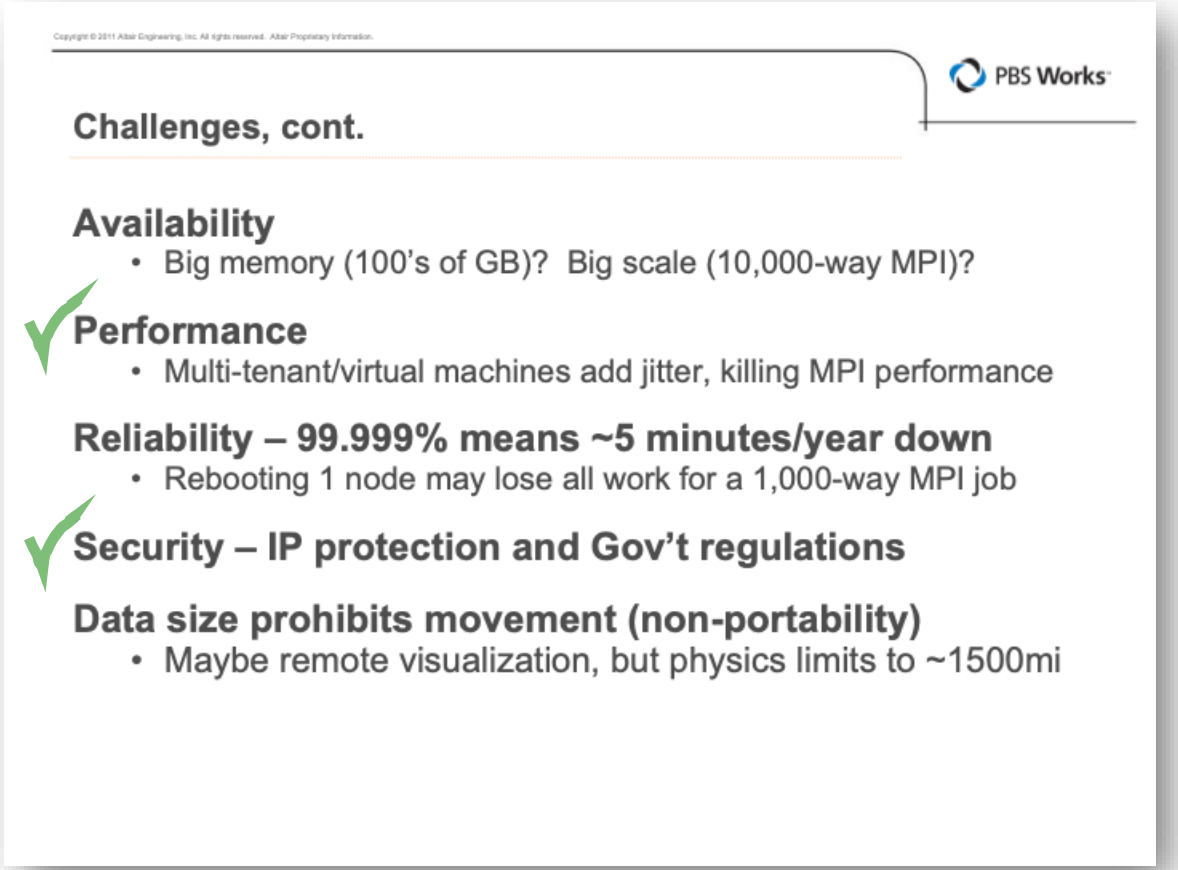

~~"Save $$$ (only pay for what's used)"~~

~~"Shrink you IT staff"~~

"Cloud has economy-of-scale"

ALTAIR

Alas, Cloud for HPC…

- Hard to get resources at scale, e.g., lots of GPUs
- More expensive (and easy to overspend) – 2-10x on-prem costs for traditional HPC sites
  - Oh, and Cloud vendors do not make it easy to control/view costs
- Different and ever changing interface requires a totally new skill set for admins
- Complex to rightsize (benchmark) given the huge set of (ever changing) choices (*)
  - Oh, and different Clouds offer different features at different values
- Prone to vendor lock-in
- Data… yeah…

(*) although, choice is generally good…

SO, WHAT WORKS FOR CLOUD FOR HPC?

Hybrid = Best Use Case for HPC Cloud

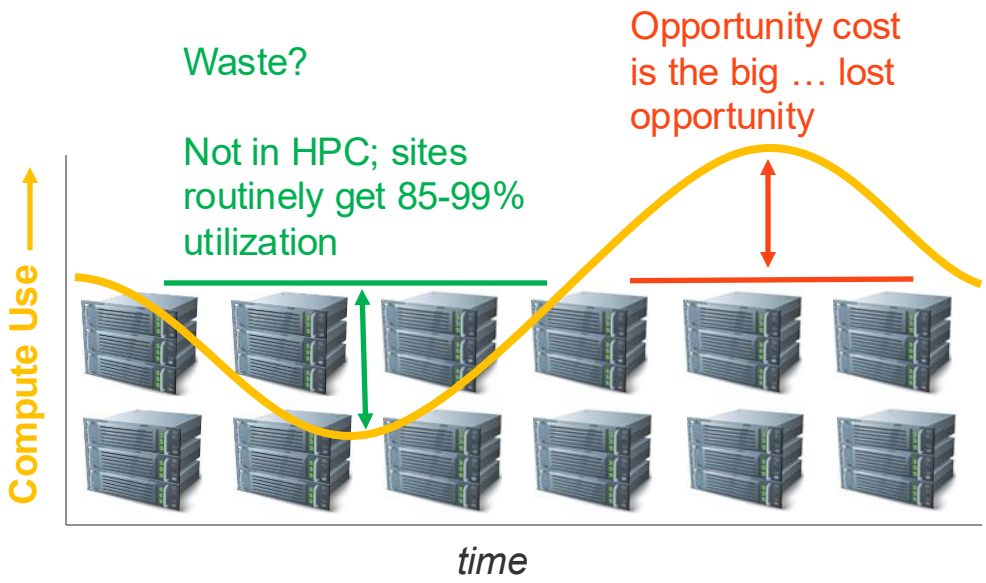

Opportunities

- Meet project deadlines
- Quickly scale for dynamic workloads
- Test new / rarely-used tech, e.g., GPUs/Quantum
- Offload “cloud-friendly” workloads
- Cover for (planned) outages, e.g., chiller
- Disaster recovery (with right contracts)

(when opportunity >> extra costs)*

ALTAIR

Note 1: Quickly Scale (On-prem vs. Cloud)

Hewlett Packard Enterprise · DELL EMC · intel · NEC · FUJITSU

Google Cloud · aws · Microsoft Azure · ORACLE

Lead Time

6 months+	Specify
	Benchmark
	Select Vendor
6 months+	Order
	Receive
	Implement
	Test
	Consume

Lead Time

Specify	1 - 6 months
Benchmark	
Select Vendor	
Request	1 – 10 mins.
Receive	
Consume	

ALTAIR

Note 2: Quickly Scale

The devil is in the details… too quickly can be inefficient…

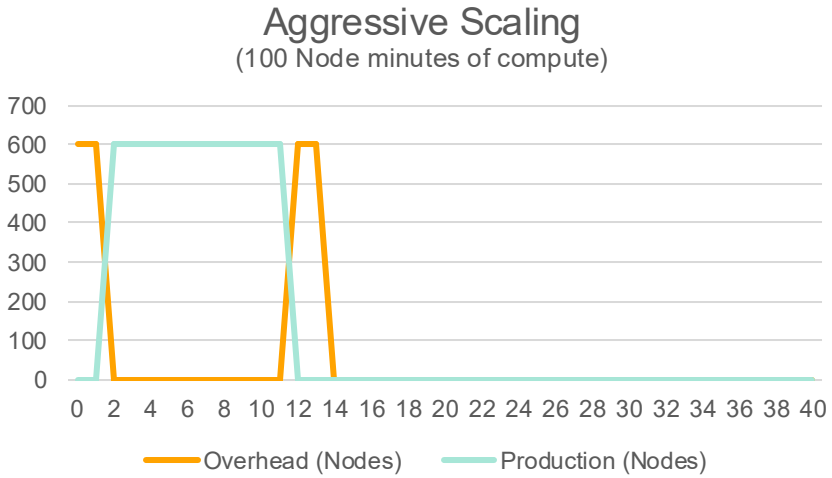

Delivery Focus

Total Cost: $5,939

Overhead: $1,697

Efficiency: 60%

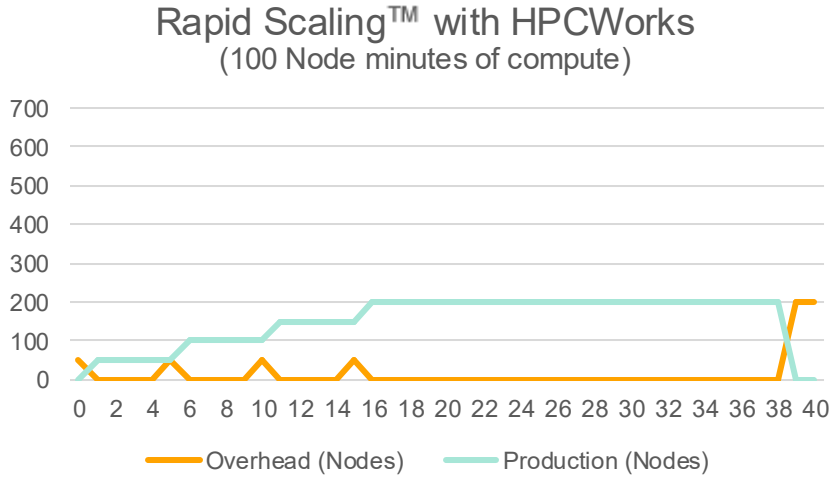

SLA Focus

Total Cost: $4,737

Overhead: $424

Efficiency: 90%

Saving: 20%

Note: In this illustrative example overhead relates to nodes during provisioning or de-provisioning which are unproductive but being paid for.

ALTAIR

Cloud Cost Control

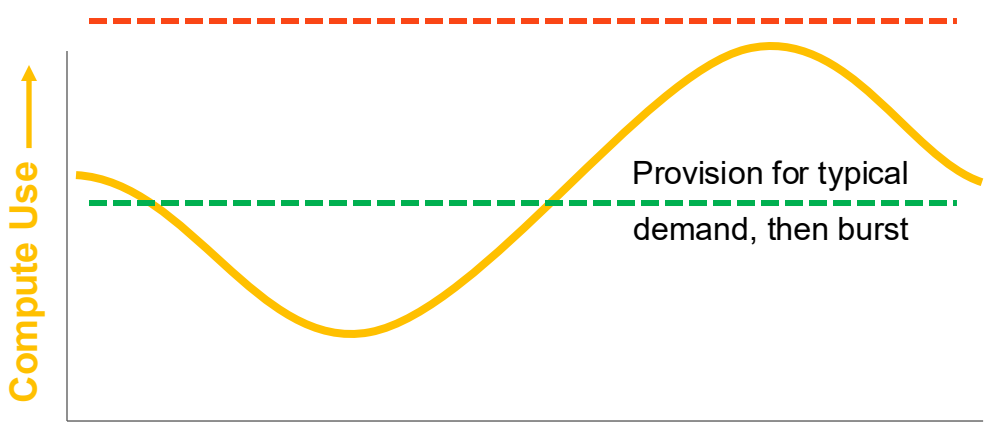

TCO
~35%

Manage and automate … everything

- Automated rules for use
- Auto-detect & shutdown idle
- Use spot, etc.
- Right-size instances
- Cost controls, quotas, et. al.
- Use WLM, e.g., PBS Pro

Percentages based on sample TCO reduction calculation based on AWS deployment example

OK, all Cloud (lift-and-shift) makes sense too sometimes, e.g., to balance CAPEX-OPEX

Getting Hybrid Cloud Right

Business-policy driven and flexible – what, why, when, where – including:

- SLAs: budget, deadlines, backlog
- Preferences: on-prem first, energy costs, …
- Filters: projects, accounts, groups, queues, priorities, software, licenses, job class, OS, node types
- Custom categories: cloud eligible, security level, database req, parallel env, …
- Fully admin controlled, but with ability for users to request exact specs, esp. for benchmarking

Open

- All Clouds, workload, WLMs, …

Automated

- Ideally, users not required to change behavior
- Scale up and scale down
- Right-size infrastructure
- Also: DevOps, deployment, tracking, …

Bulletproof accountability and full visibility

- Cost: estimation, real-time use, quotas, limits, reporting
- End-to-end logging, for fast troubleshooting and to keep CSPs honest

ALTAIR

EXAMPLES

Altair Cloud: Subaru

Efficient Crash Simulation

Tier 1 Automotive Manufacturer

- Oracle Cloud Infrastructure (OCI)
- Apps requiring Infiniband:
  OCI bare metal instances
- 16 Instances/Job, circa 300 max.
- Bursting from a single cloud native PBS Pro cluster

Results

- **Better time-to-market**
- **Bare-metal speed**

Altair HPCWorks

Powering Drug Discovery

Johnson & Johnson – Janssen Pharmaceuticals

- Drug discovery – COVID 19 vaccine
- Capacity @ scale (3x bigger)
- HPC in Cloud, ~20 clusters, all different
- Multi-tenant environment
- ~1.500 users; moving to 10K cores

Results

- **Scale (fast) to match demand**
- **On-demand science**

> The solution including Altair Grid Engine and NavOps enabled us to scale in the cloud and grow our HPC infrastructure to support critical science and research, including COVID-19 vaccine development.
>
> Martin Dellwo, Manager, Advanced Computing, Janssen R&D

Altair HPCWorks

Energy efficient powertrains

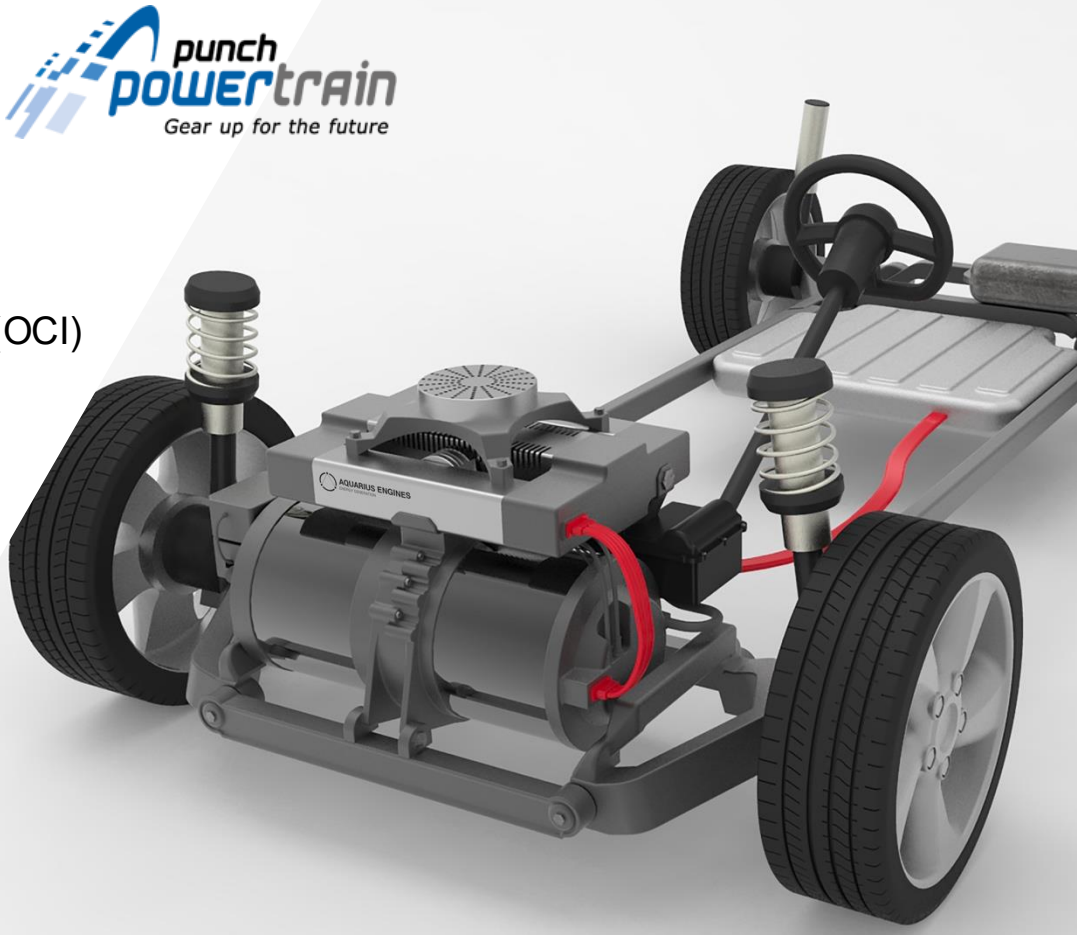

Punch Powertrain

- Multi-cloud: Oracle Cloud Infrastructure (OCI) and Google Cloud Platform (GCP)
- Apps requiring Infiniband: OCI bare metal instances
- GCP for other apps
- Bursting from a single *PBS Pro* on-prem cluster into both clouds

Results

- **Route work for best ROI**
- **On-prem & multi-cloud**

Altair HPCWorks

Harnessing the power of wind, sun and water

GE Wind Energy

- AWS Spot Instances in multiple AZs within the local region
- Round robin on AZs to harvest spot instances; spot pricing support
- Mix of Windows and Linux
- Batches of 5K job arrays
  500+ burst nodes

Results

- **Fast, cost-effective scaling**

Altair HPCWorks (Beta/PoC)

NCAR

AMPS: Predicting Antarctic weather

National Center for Atmospheric Research

- Coverage for Maintenance Windows
- Provide Mission Critical Forecast Data
- Automated Deployment, Configuration and Operations
- Scale 100s of IB Enabled Instance (AWS EFA)
- Automation of Lustre FSX Sync-ed to S3 Bucket
- Supporting Antarctic Team Safety

Results

- **Reliable forecasting (thru outages)**

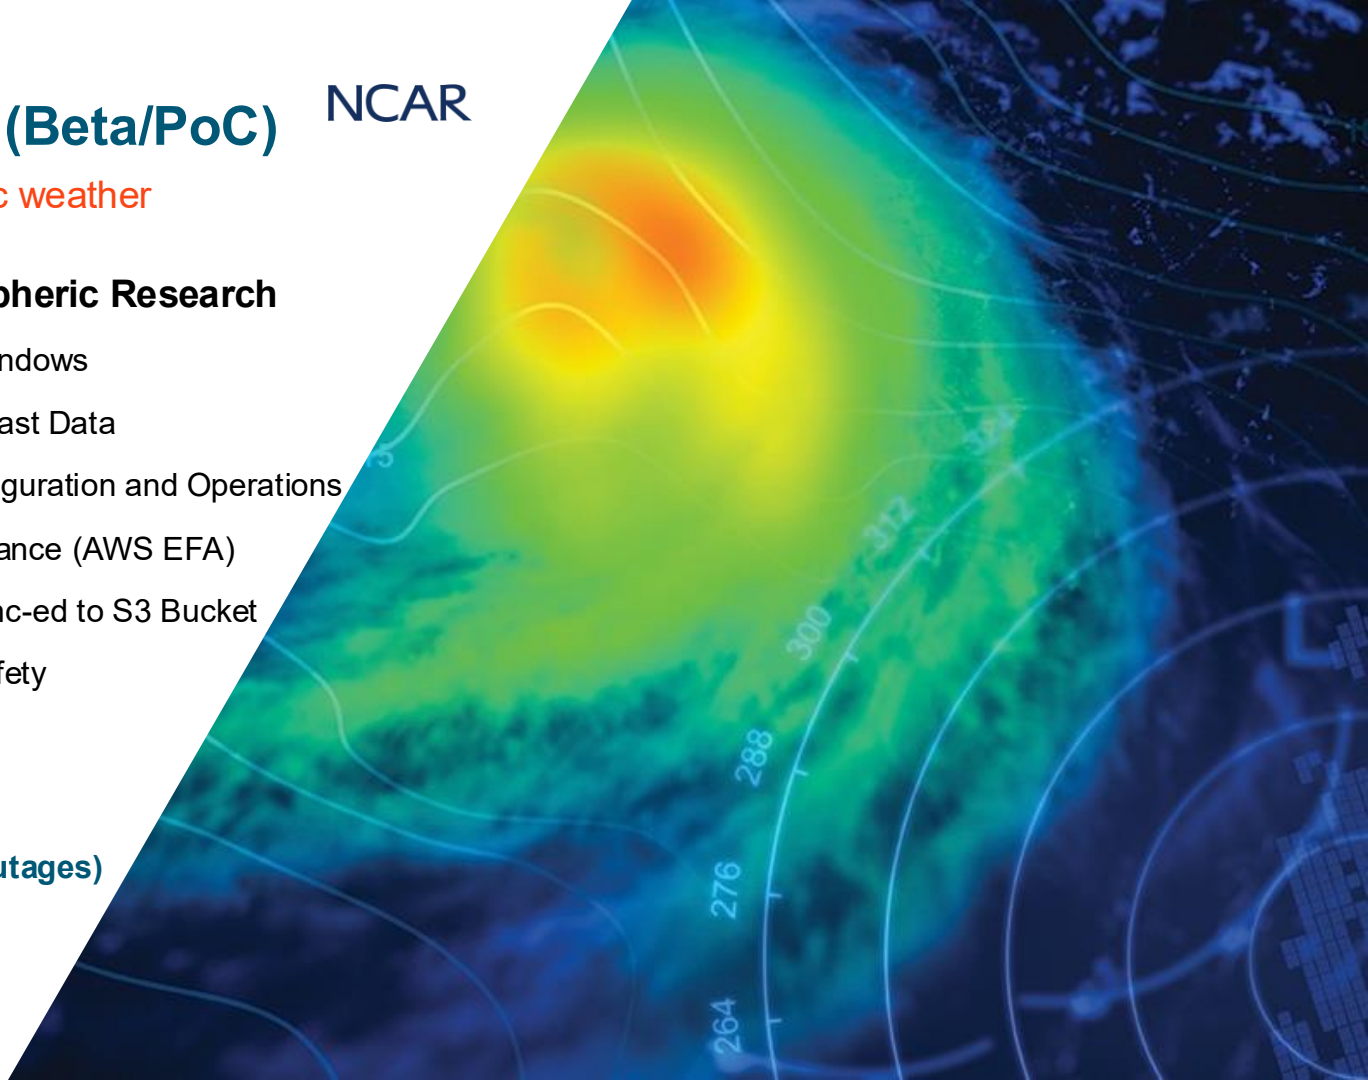

ALTAIR NAVOPS®

ENTERPRISE CLOUD SCALING FOR HPC

ALTAIR

Altair NavOps

Any Workload, Any Cloud, Any Scale, Anywhere

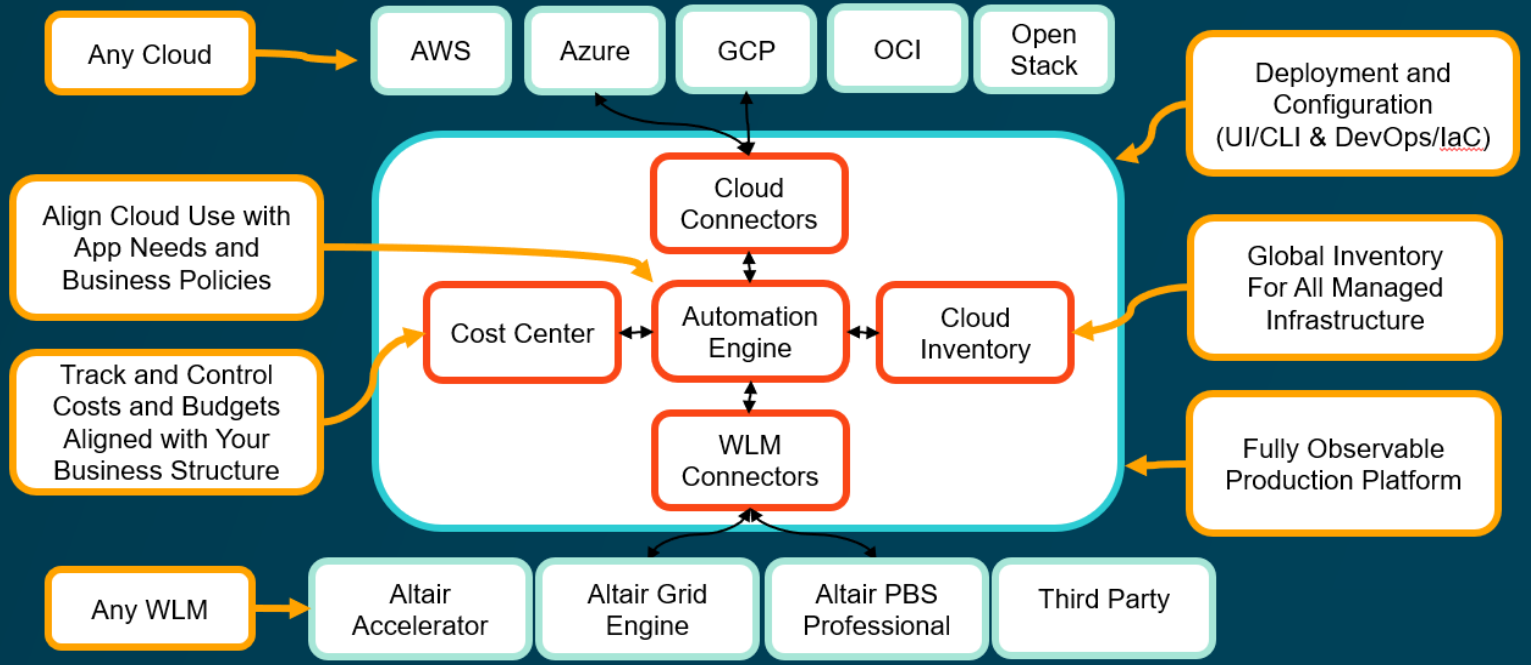

ALTAIR

Altair NavOps

Hybrid Multi-cloud scaling & bursting

- All uses – hybrid, Cloud native, lift-and-shift
- All workloads – AI & HPC (PBS Pro, AGE, Accel., Slurm, K8s, …)
- All clouds -- no lock-in, exploit arbitrage, inc. specific capabilities, e.g., spot
- All scales –naïve user to global enterprises

Enterprise-level DevOps/IaC platform

- Automation Engine with Rapid Scaling™ delivers business needs fast and efficiently
- API-first (w/great GUI), budgets, limits, cluster deployment, full observability

Centralized cost controls and dashboarding

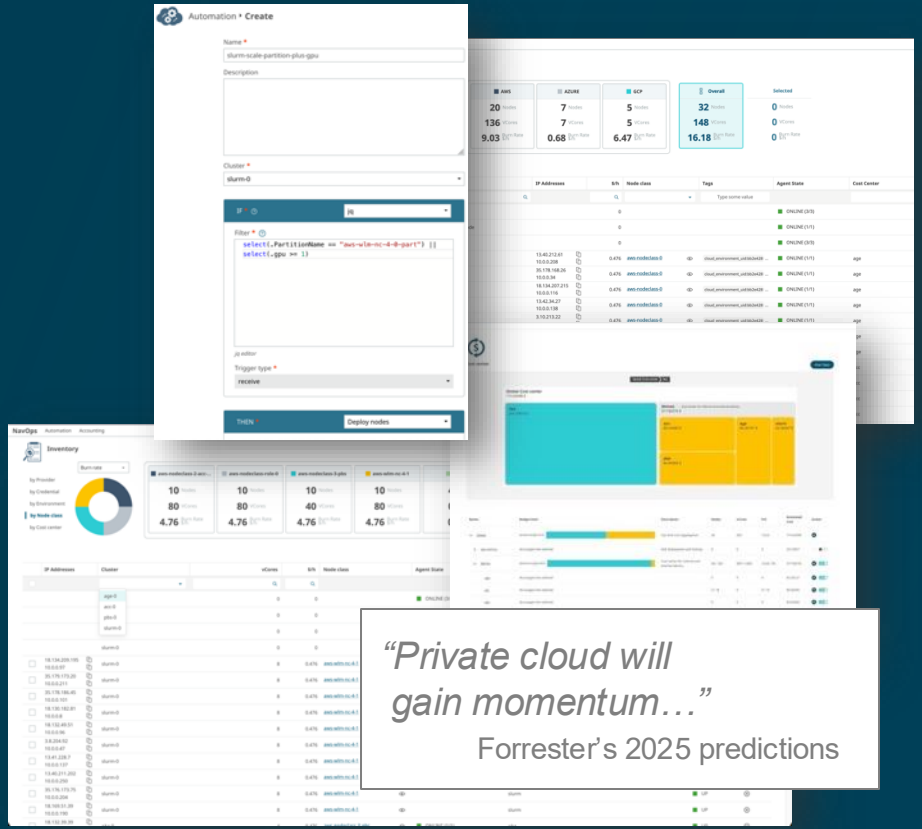

ALTAIR

Altair InsightPro™

Manage to business priorities: 360° visibility

- Projects, costs, budgets, quotas
- Rightsize infrastructure, eliminate waste
- Chargeback, tune, troubleshoot

Analytics for HPC Infrastructure and Use

- Turnkey → custom expansion → open APIs
- Data for “everything”: workload, hardware, software licenses, performance, faults, $/costs, energy, carbon, …

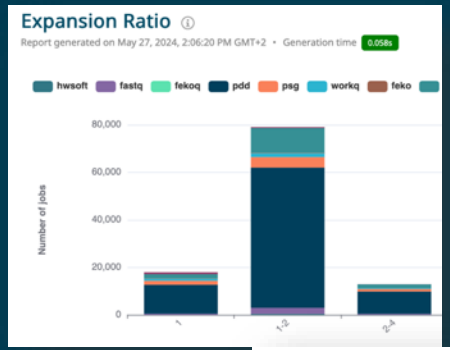

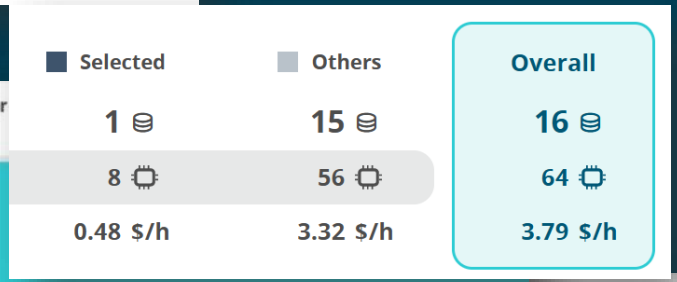

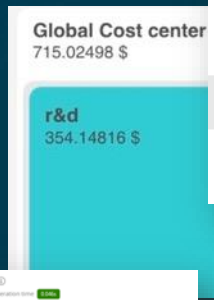

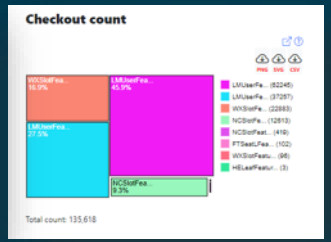

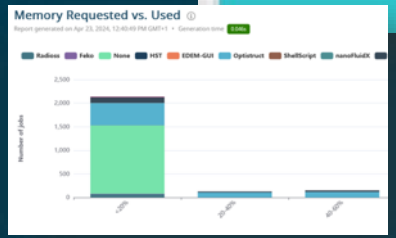

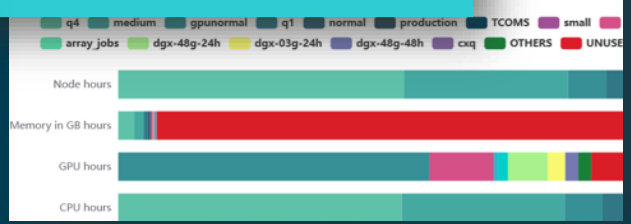

Altair NavOps: Example

Scaling for SLURM

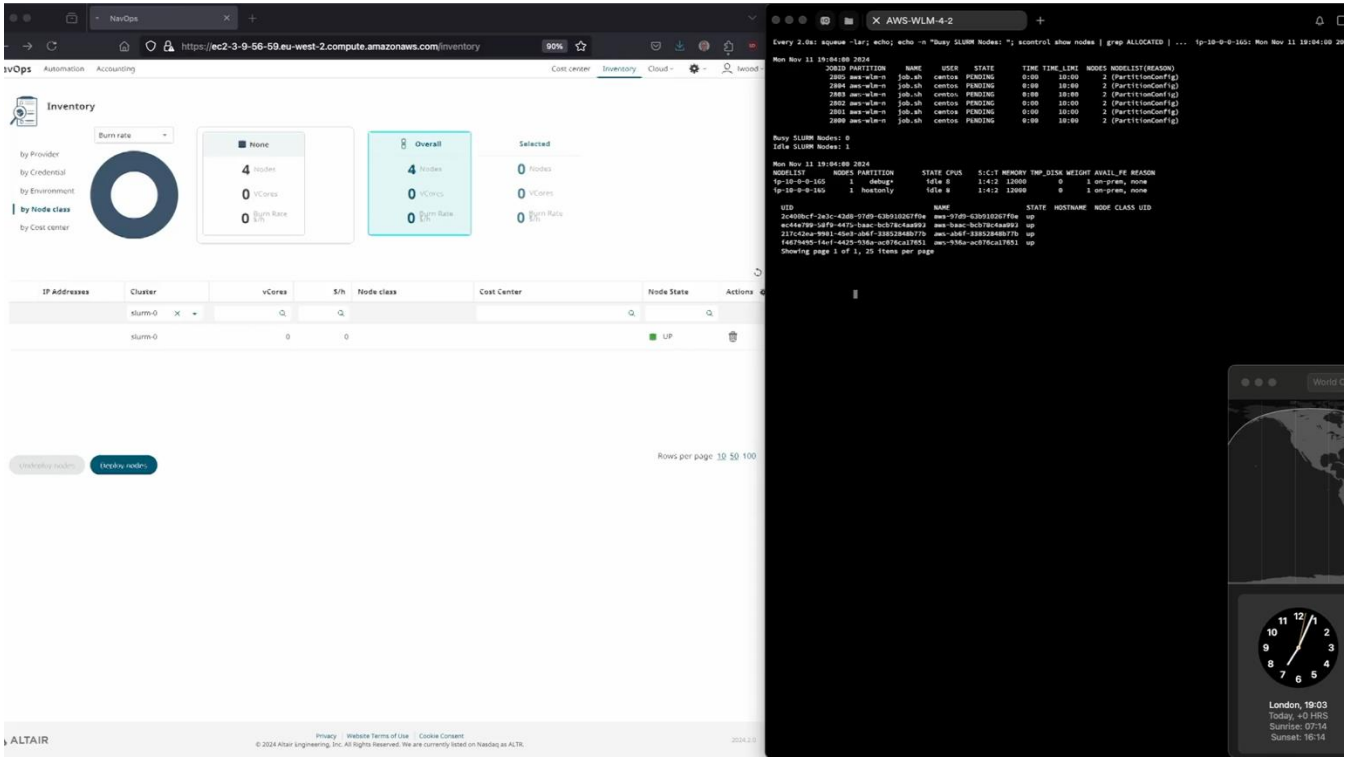

ALTAIR


Get the right resources
to the right person (agent)
at the right cost
and at the right moment.
Powered by Altair® HPCWorks® and Altair® RapidMiner®
ALTAIR