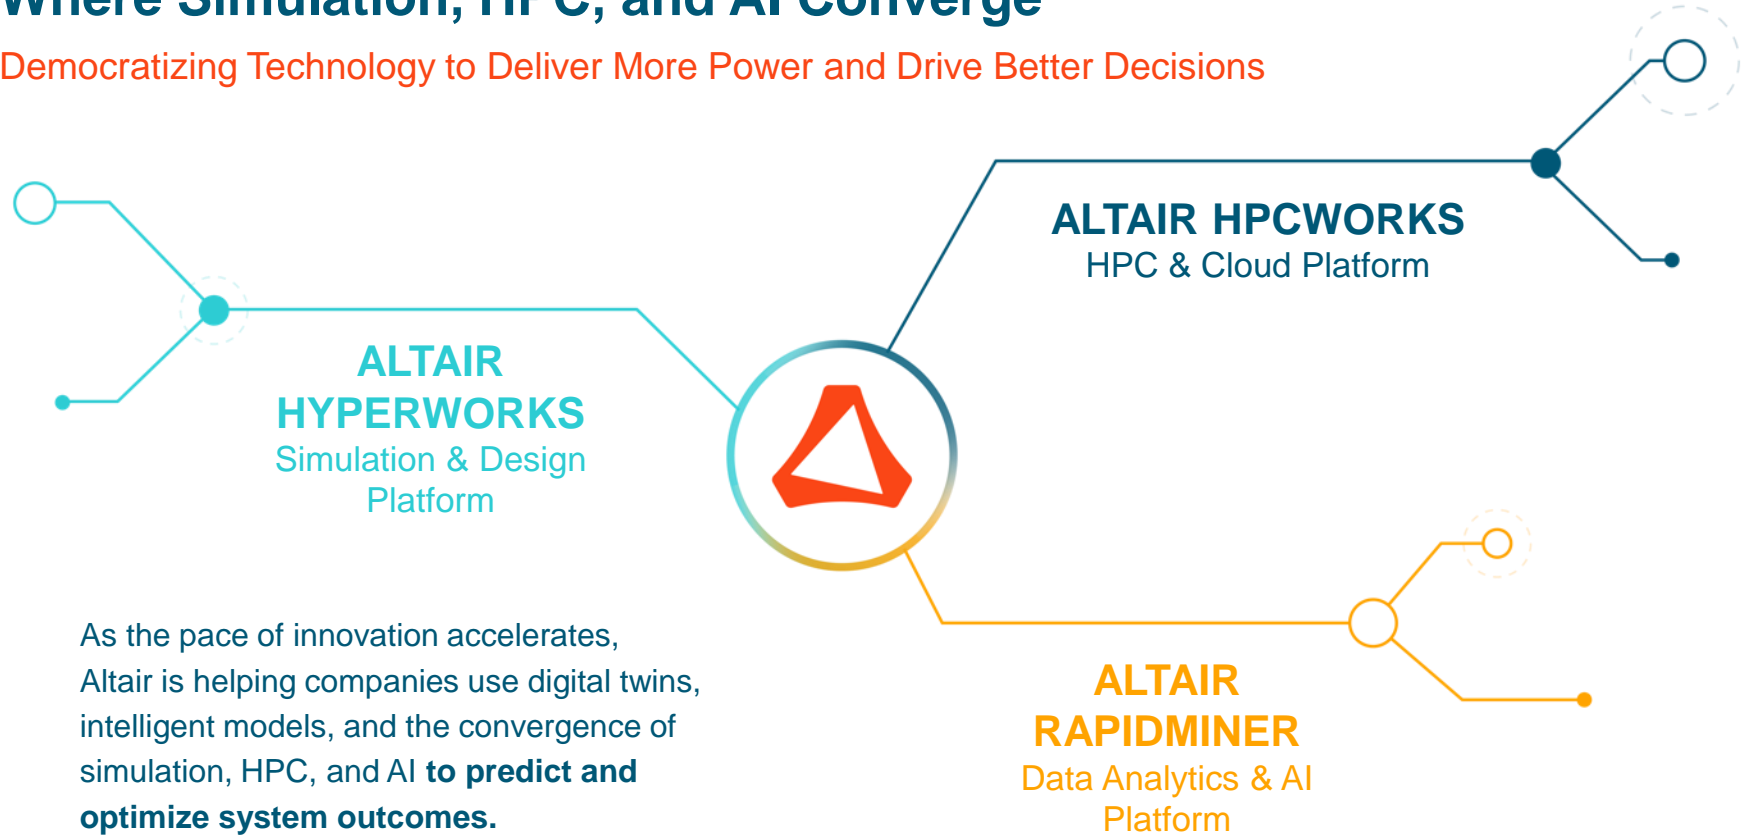


**ALTAIR® LIQUID SCHEDULING®:
THE BIGGEST CHANGE TO HPC IN 30 YEARS**

Dr. Rosemary Francis, Chief Scientist HPC

Where Simulation, HPC, and AI Converge

Democratizing Technology to Deliver More Power and Drive Better Decisions



As the pace of innovation accelerates, Altair is helping companies use digital twins, intelligent models, and the convergence of simulation, HPC, and AI **to predict and optimize system outcomes.**

AI workloads are taking over

Data analytics

ML on place of solvers

Large language models

Self driving cars

Edge, IoT, streaming data

Hybrid HPC - AI workflows



How are AI jobs different from HPC

Training vs inference

Run at the edge

Often High Throughput, not High Performance

Often Bursty

Often have tight deadlines for reasonable ROI

Platforms are becoming more diverse

Cloud


Hybrid Cloud

Kubernetes

Edge




Hardware is getting more complicated




CPU



GPU



DPU



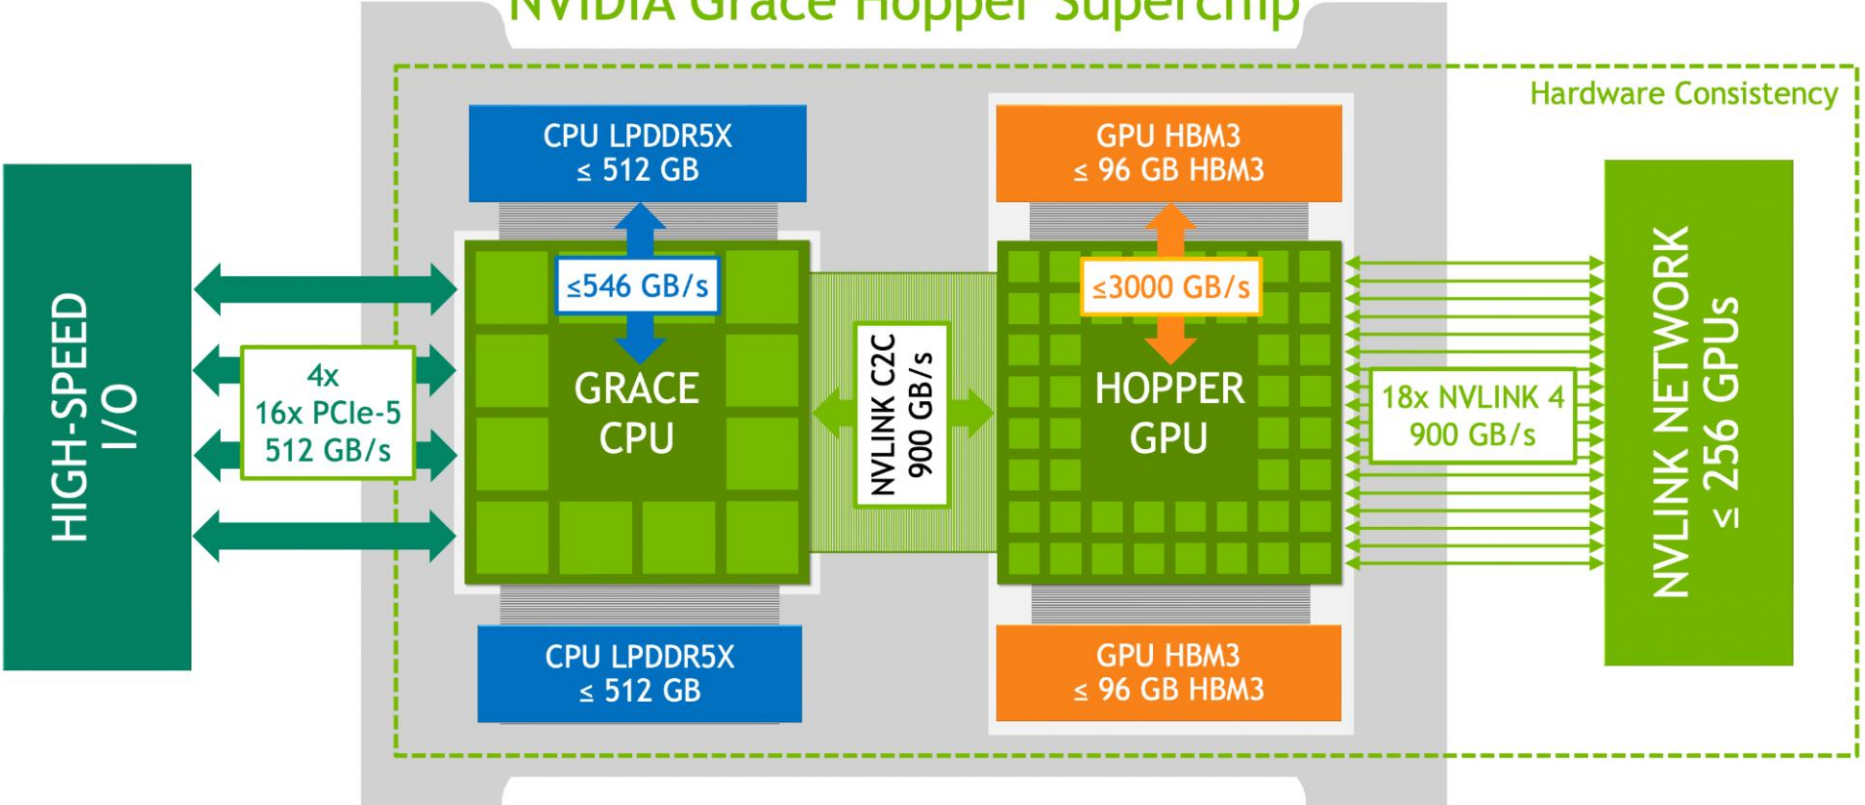
Networking



Storage

GPU SoC mega chips


NVIDIA Grace Hopper Superchip




NVLINK NETWORK
≤ 256 GPUs

GPU SoC mega chips

AMD CDNA™ 2
Coherent Memory Architecture



AMD CDNA™ 3
Unified Memory APU Architecture



Inference and AI chips

Groq



Graphcore

Saliency

Cerebras

SambaNova

Habana



Scaling challenges

End to Moore's law

Cores increase but node count does not

Exascale

Distributed HPC

Cloud is not infinite

Challenges as HPC increases in complexity

HPC today covers a diverse range of applications, technologies, and hardware

- High overheads in integration and support

Users have to choose where to run their codes

- Leading to low utilization

Administrators have low visibility and control over resources

- Causing poor return on investment (ROI)



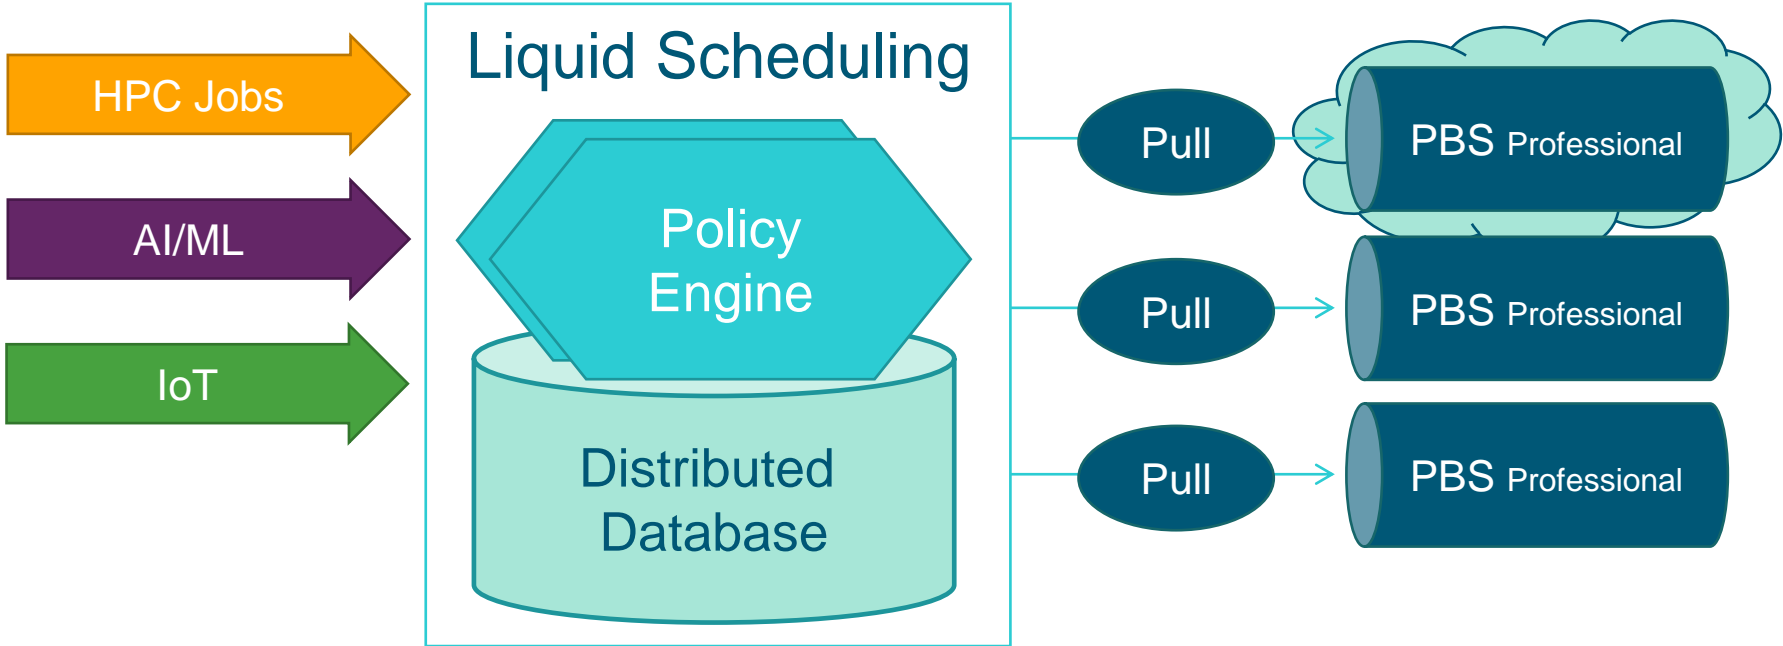
Liquid Scheduling

Liquid Scheduling creates an HPC platform that delivers scalability for the next generation of supercomputers.

- **Leap in performance** and scalability
- **Support diverse new workflows** and exotic compute
- **Break down silos** by connecting multiple HPC clusters
 - Improve utilization and access to resources
- **Built using a modern web-scale technology stack**

Liquid Scheduling

Bringing together new workloads and compute platforms



Key architectural concepts

Multiple workload managers queue jobs only when they are about to run them

Jobs execute on the first available resources

Liquid Scheduling applies global policies

Global control over fairshare and prioritization for better utilization

Submission through PBS Professional CLIs

No change for end users or applications

Built for future scalability needs

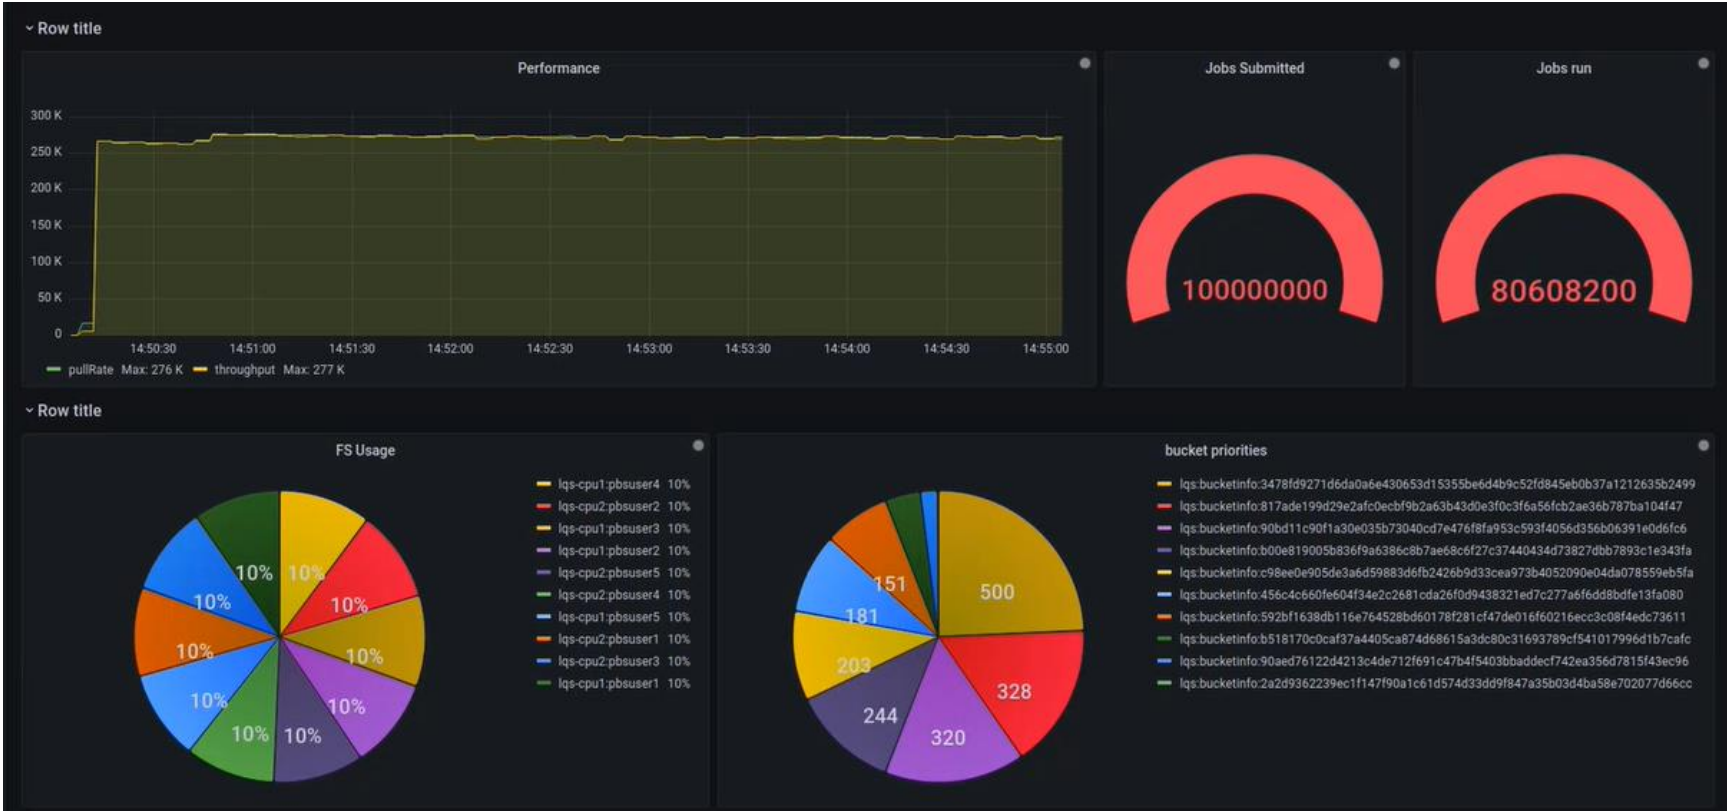
Using modern, web-scale technologies and streaming architectures

What does this mean to PBS Professional customers?

- Schedule workloads automatically across multiple PBS Professional clusters
- Fairshare implemented across multiple PBS Professional clusters
- Resilience against qstat overload
- Better performance and scalability even when using just one PBS Professional cluster

Demonstrating scalability

100M jobs across eight PBS Professional clusters with fairshare delivering >250K jobs per second



Why is Liquid Scheduling **so fast**?

- Separation of scheduling policies from resource management
 - No bottleneck component that needs to know everything
 - Policies are applied to groups of similar jobs
 - Workload managers only handle small queues of jobs that are ready to run
 - Even with just one PBS Professional cluster, performance is improved

What makes Liquid Scheduling so disruptive?

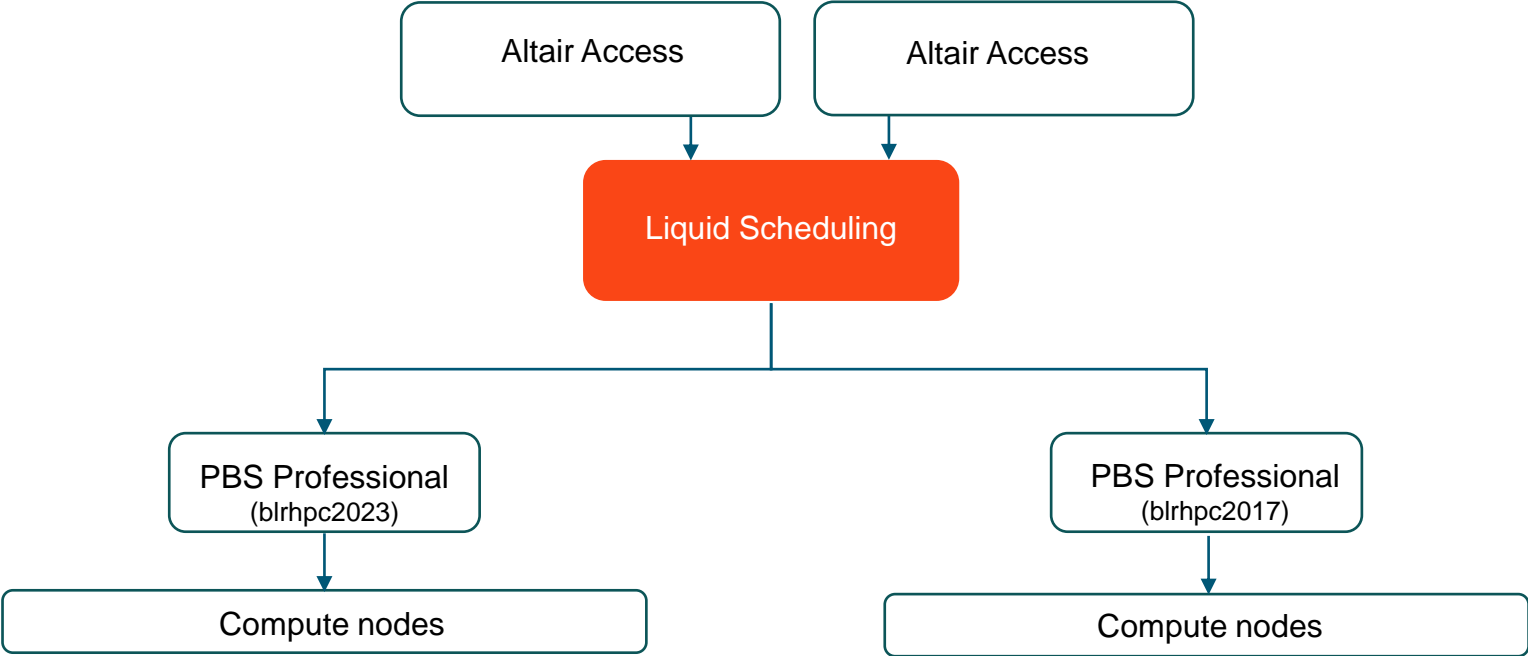
Altair is the world leader in ensuring workloads run in the right place at the right time.

- ✓ Connecting distributed compute resources
- ✓ Users no longer have to choose where to run their jobs
- ✓ Fairshare correctly applied over all compute resources
- ✓ Overall system utilization is improved
- ✓ System monitoring and accounting is now unified

Liquid Scheduling is already in production at Altair

- Liquid Scheduling has been deployed to unite some of our PDD clusters
- Integration with Altair® Access™ to leverage the whole Altair® HyperWorks® stack
- Users and applications migrated one at a time without disruption
- **Feedback – jobs run sooner, otherwise the system looks the same**


Liquid Scheduling on our PDD clusters



Solution: Managing multiple workload managers

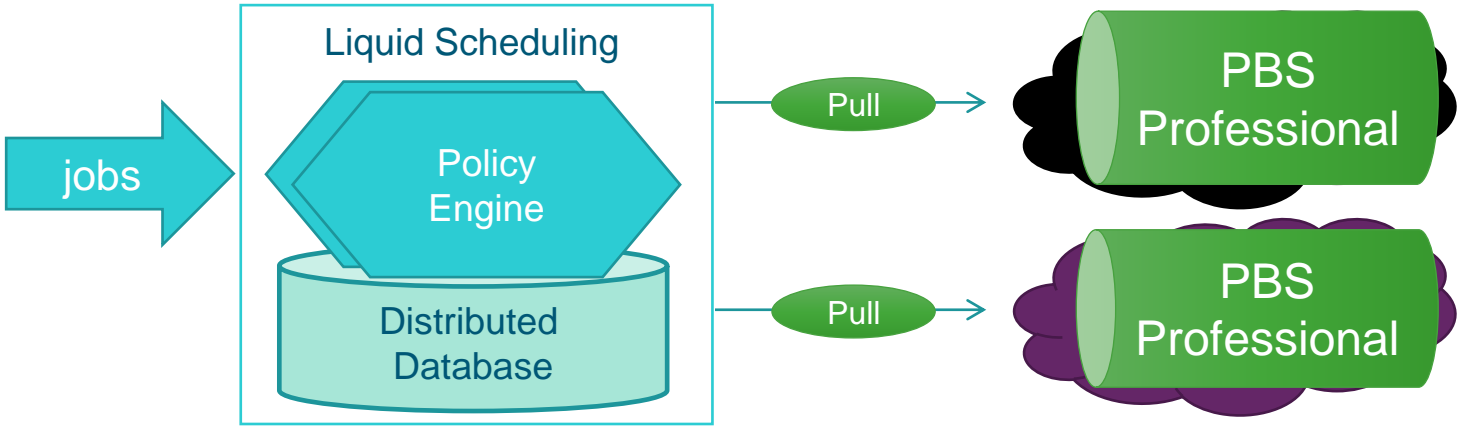
- ✓ Users no longer have to choose where to run their jobs
- ✓ Fairshare and quotas apply to all compute pools and are centrally managed
- ✓ Overall system utilization is improved
- ✓ System monitoring and accounting is now unified

Submission through CLIs
and APIs of our WLMs
→ no change for end
users



Example: Scaling in the cloud

- ❖ Workloads are bursty, very parallel, and business-critical
- ❖ Public cloud can't always deliver the scale needed
- ✓ Deploy Liquid Scheduling to distribute the workload across multiple cloud vendors



Altair® Liquid Scheduling™

- **Makes Altair® PBS Professional® faster and more resilient**, even with only one cluster
- **Improves PBS Professional scaling** by breaking down silos between multiple clusters
- **Users and applications don't need to change anything** to have their jobs run sooner#

Architected for the future of HPC