



Hewlett Packard
Enterprise

HPC and AI in hybrid cloud environments with HPE GreenLake



César Gómez – WW HPC Solutions
Architect

Cáceres (Spain), September 14, 2022



Looking at the Gordon Bell Prize


- 1 GFlop/s; 1988; Cray Y-MP; 8 Processors
 - Static finite element analysis
- 1 TFlop/s; 1998; Cray T3E; 1024 Processors
 - Modeling of metallic magnet atoms, using a variation of the locally self-consistent multiple scattering method.
- 1 PFlop/s; 2008; Cray XT5; 1.5×10^5 Processors
 - Superconductive materials
- 1 EFlop/s; ~2018; ?; 1×10^7 Processors (10^9 threads)



Jack Dongarra

Mérida, 27 Abril, 2010


CONSIDERATIONS FOR WORKLOAD PLACEMENT



¹ Source: Hyperion research 2022


HPE GREENLAKE

EDGE-TO-CLOUD PLATFORM



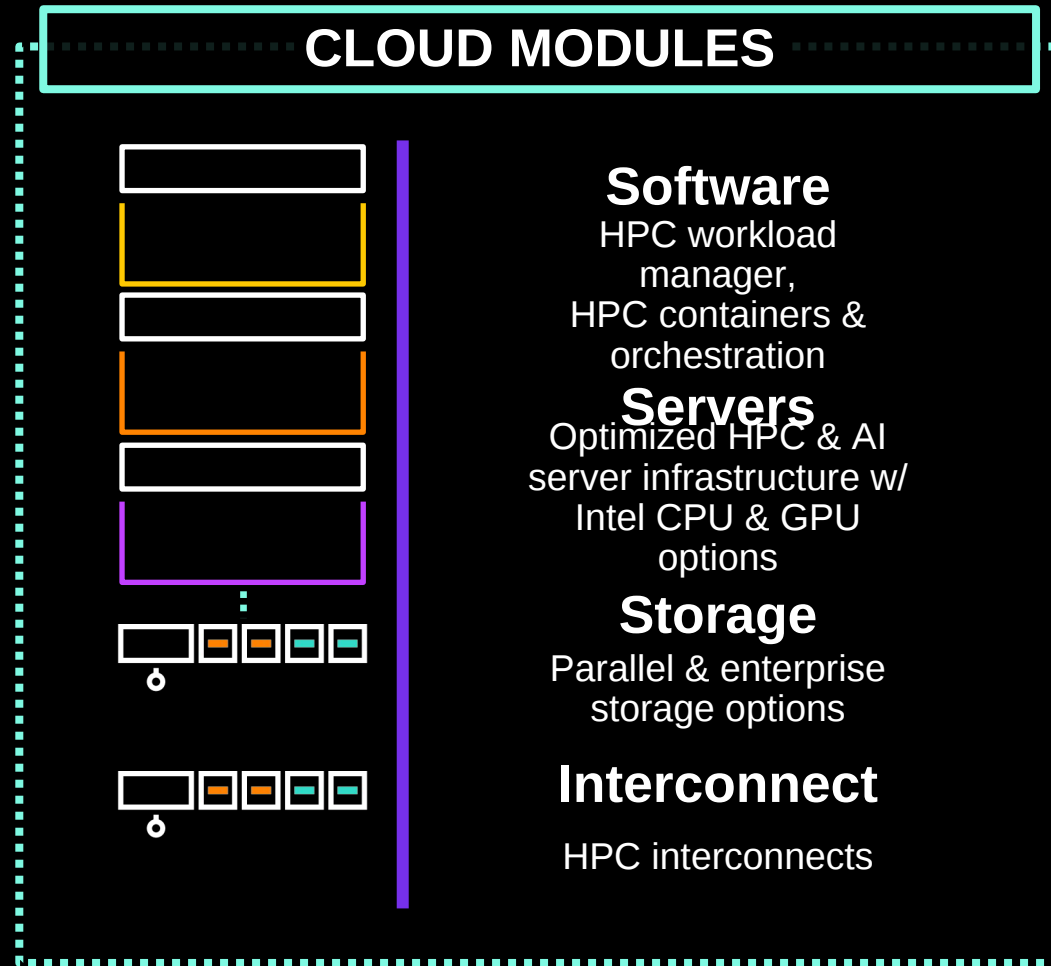
THE CLOUD THAT COMES TO YOU

¹ Reserve may apply



HPE GREENLAKE FOR HPC

Simplicity and speed with cloud modules



Mix and match compute
and storage modules

Validated and
optimized for
performance

Supports multiple
HPC/AI/ML workloads

NEW HPE GREENLAKE FOR HIGH PERFORMANCE COMPUTING

- HPE Cray and Intel technologies delivering throughput unprecedented in a cloud service

Extensive HPC Partner ecosystem of value-added software and services integrate with HPE GreenLake for HPC by a partner program, with carefully design business interfaces

HPE Apollo 2000/6500 Systems

More flexibility to open up AI, Machine Learning and more HPC techniques

GPU Enhancement

Support for NVLINK, NVIDIA A100, A40, A30 in increments of 2-4-8 accelerators

HPE Slingshot

Cray networking technology with extremely high speed, tunable Ethernet-based interconnection supercharging performance

HPE Parallel File System Storage

Scalable, high performance storage system that can match the other components, delivering unprecedented throughput

Lower entry point

Reduce risk of introducing HPC, test workloads with HPE GreenLake for HPC and scale as needed

WHY IT MATTERS

- Hybrid Options

Flexible Hybrid models for customers, offer elasticity of their HPE GreenLake for HPC service

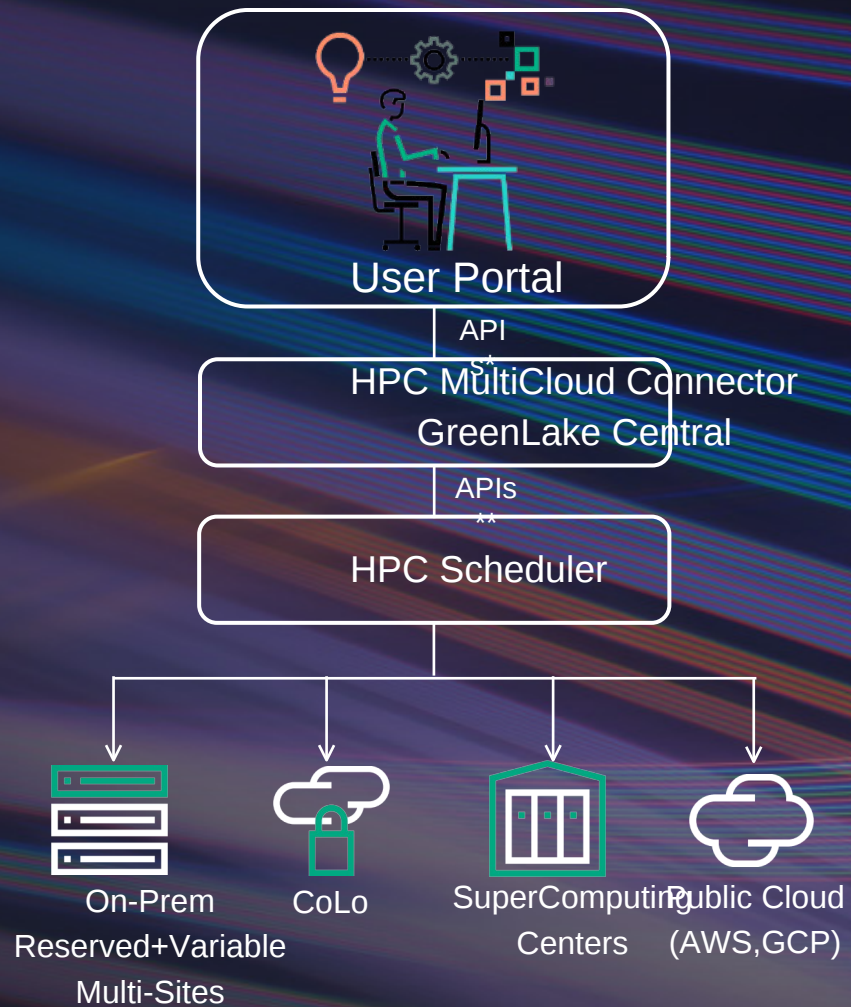
HPE GreenLake for HPC-to-private cloud based on HPE GreenLake for HPC or to-a public cloud

Multi-cloud connector APIs (Hybrid Cloud APIs) that we will publish and drive to become industry standard on how to program submitting HPC jobs to a diverse pool of computing

Ability to orchestrate data-center scale workflows with user-defined policies to determine best computing target where to execute a job

Multi-Cloud Connector

Performing hybrid capability to submit jobs to an external cluster



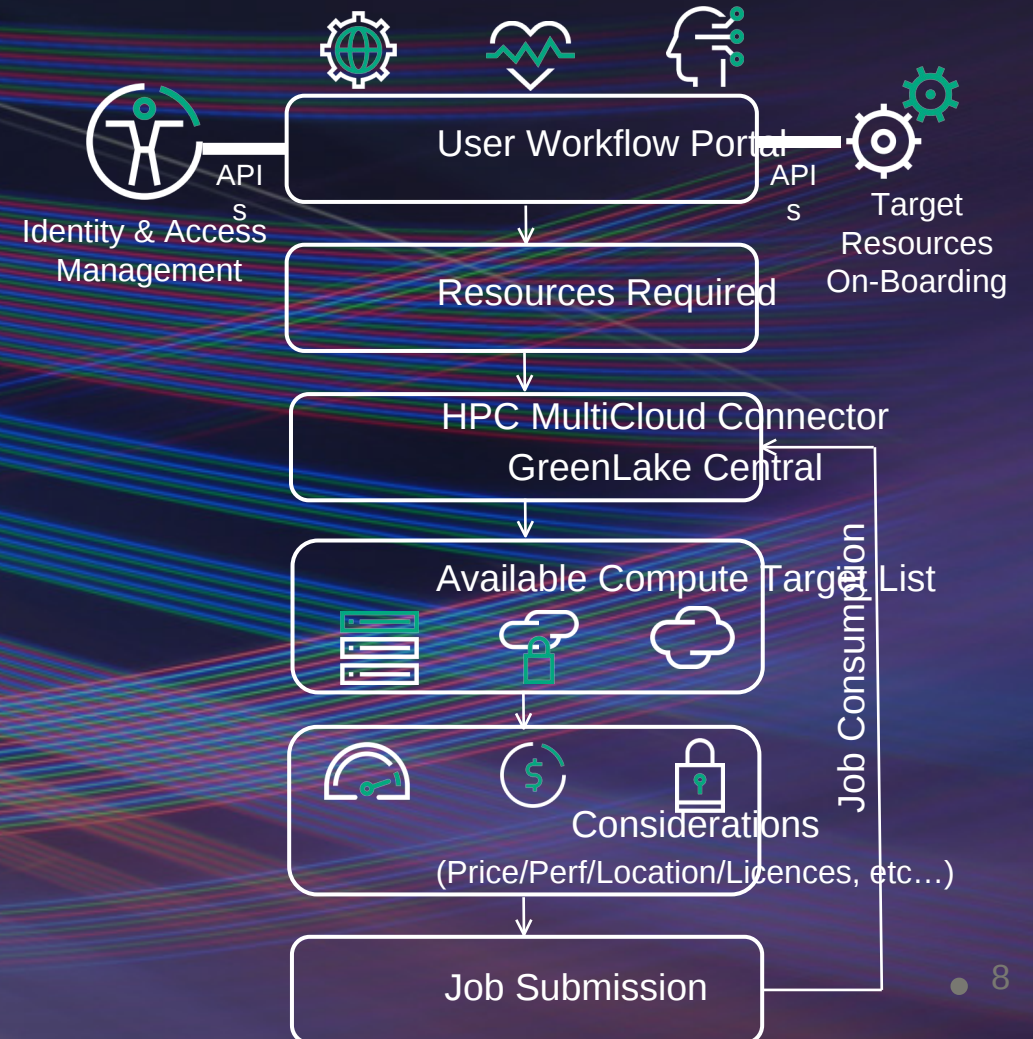
* APIs: job submission, resource filtering

** APIs: on-boarding, provisioning

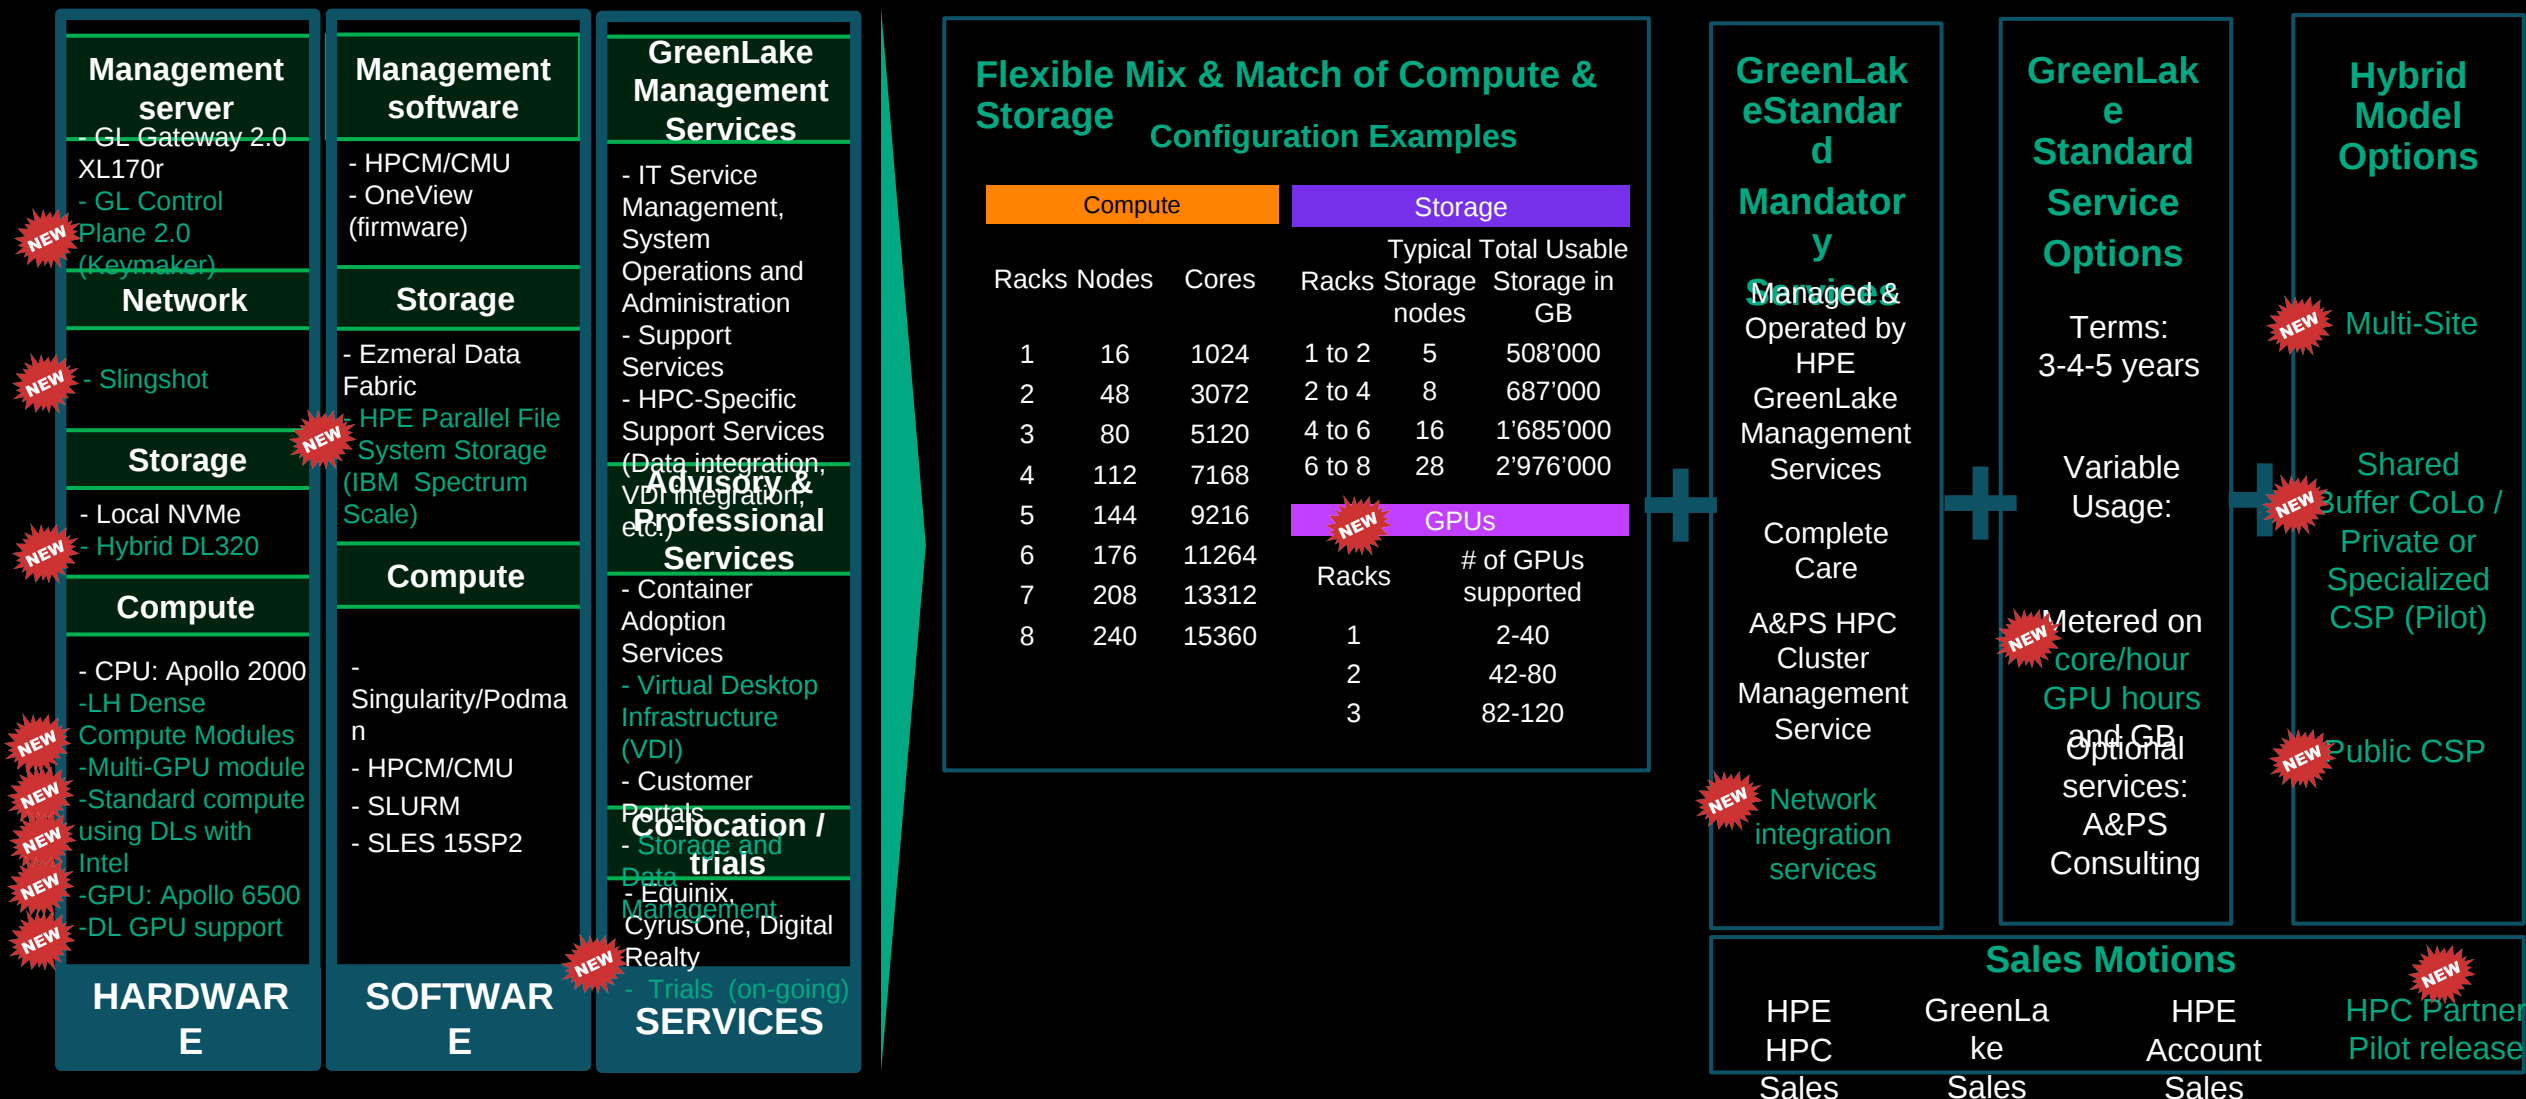
Workflow-aaS

Applied to different verticals:

- Manufacturing-aaS (d3View)
- HC and LS-aaS (3D structures of biomacromolecules)
- Large scale training-aaS (SmartSim/Determined.ai)



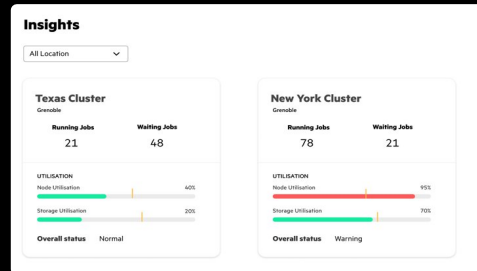
WHAT IS UNDER THE HOOD?



AN END-TO-END HPC AND AI EXPERIENCE

HPE GreenLake platform changes the experience

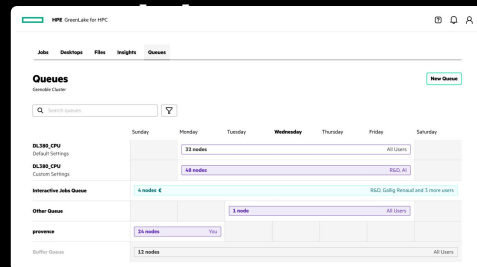
Configure the cluster



HPC
ADMIN

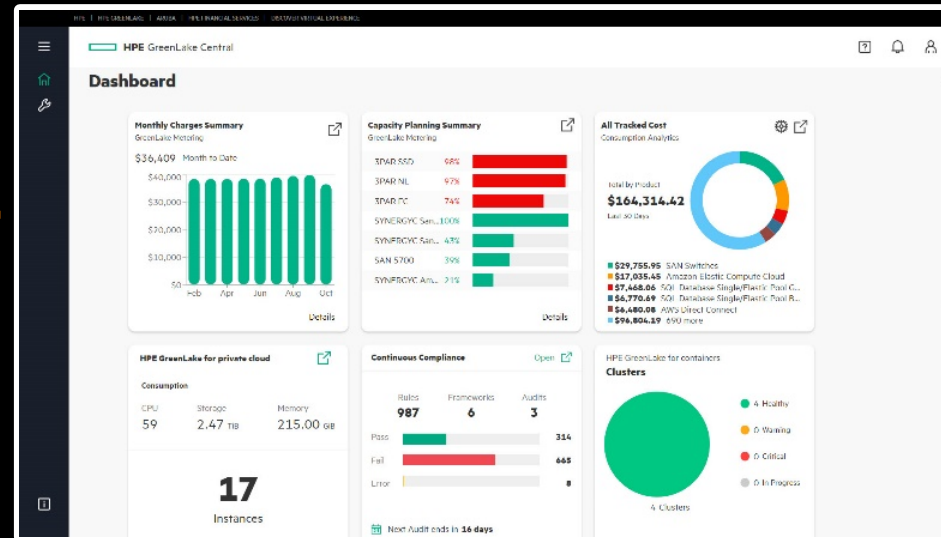
HPC
ADMIN

Launch &
monitor HPC



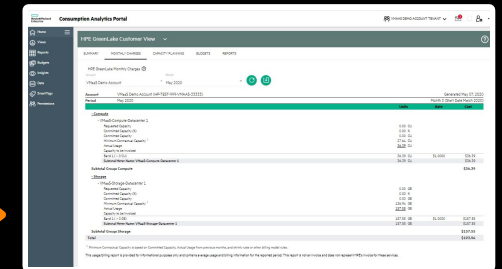
Meter costs
and usage

HPE GreenLake Central

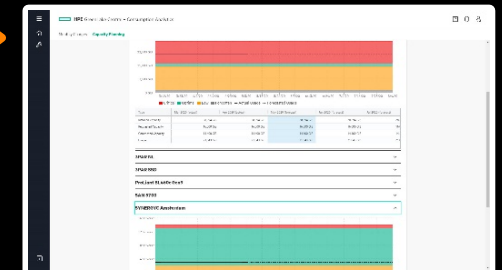


CIO/LOB

View costs and usage



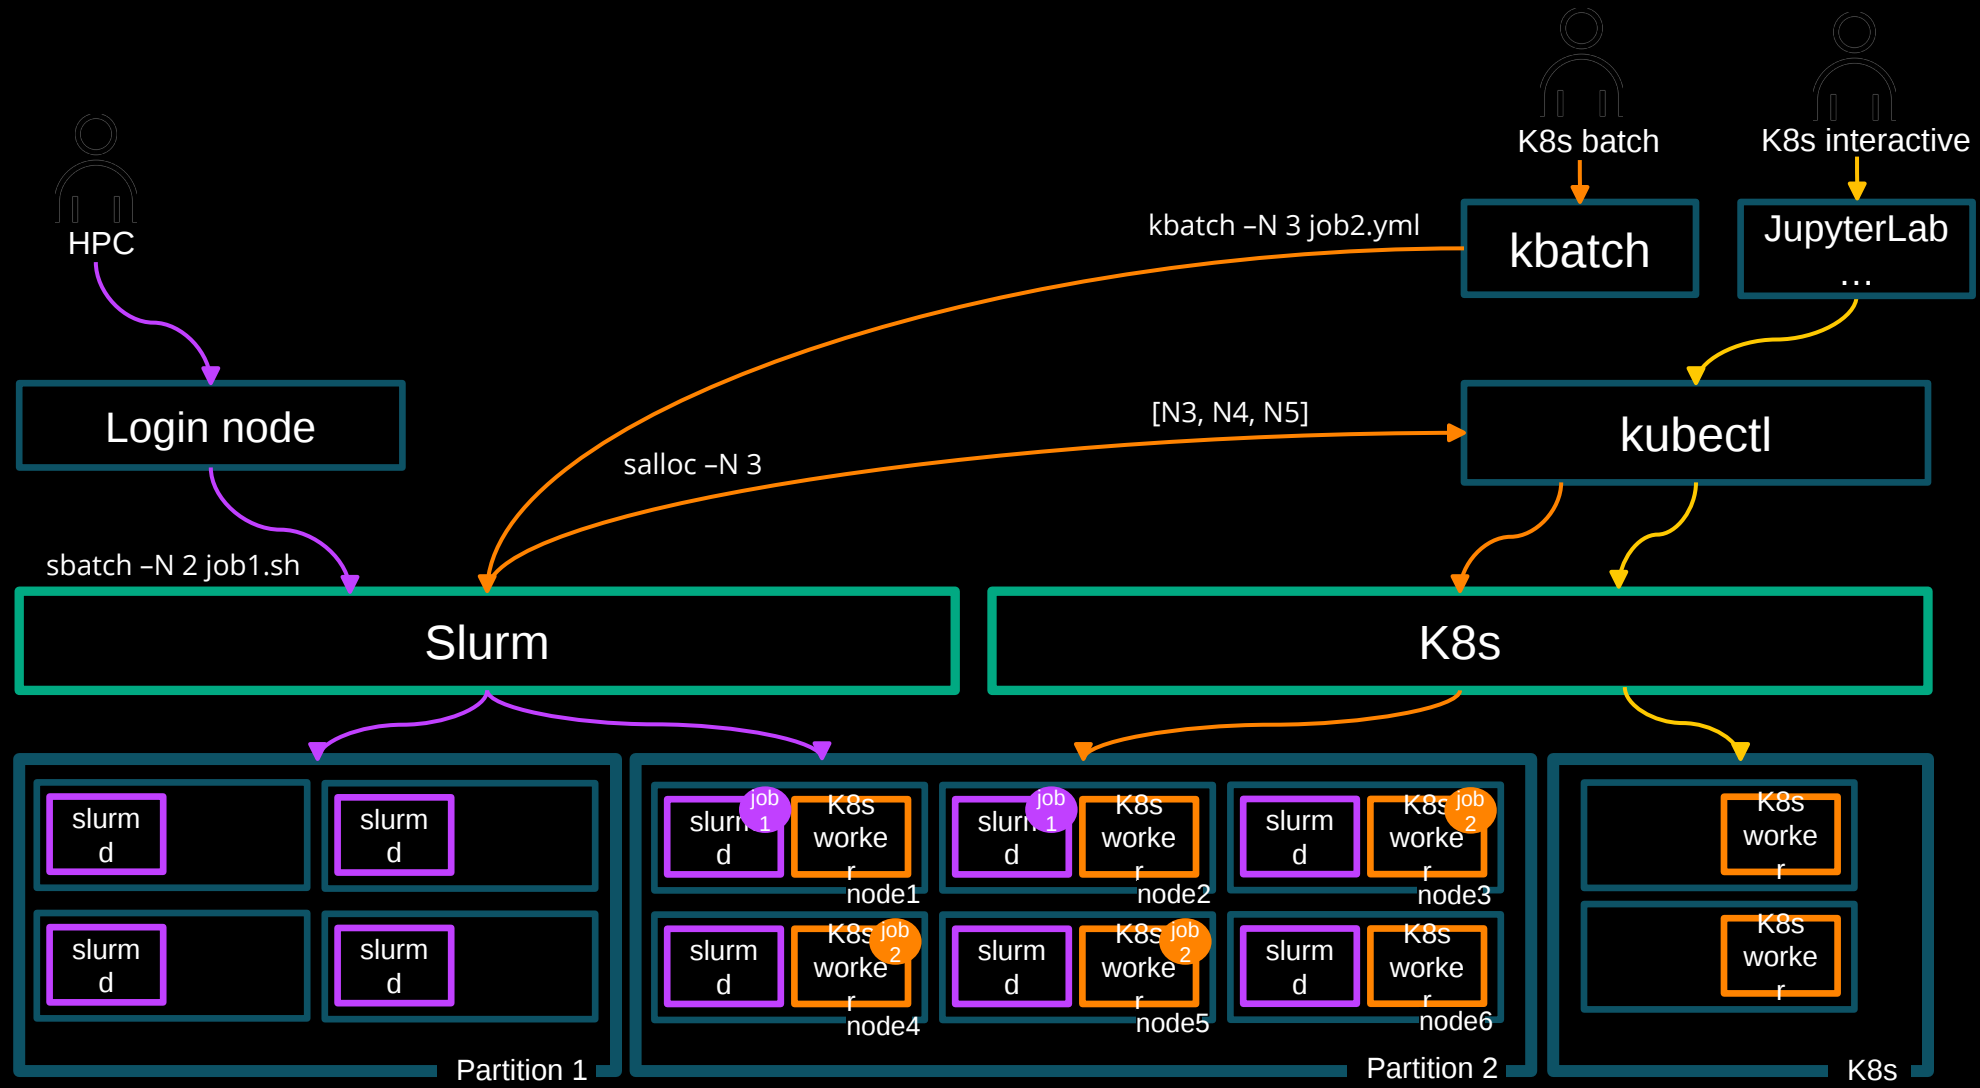
Forecast capacity



WHY IS CONVERGENCE OF HPC, AI, AND CLOUD IMPORTANT?

- To take advantage of cloud technologies for HPC/AI-native infrastructure
- To use large scale infrastructure for HPC and AI workloads
- To provide on-demand computing capacity without changing the software stack
- To have tailored or highly configurable software and workload deployments
- To abstract the complexity of running the same application in different HW and SW infrastructures
- To have the ability to run user-defined software stacks using containers

CONVERGED SLURM/KUBERNETES CLUSTER



VISION: BASIC TO EXASCALE

	MVP	Expand	Exascale
HPE GreenLake Services capabilities	<ul style="list-style-type: none"> • HPC GreenLake Central integration • Cost analytics • Self Service interface Manage clusters and batch queue • Optional A&PS Services available 	<ul style="list-style-type: none"> • End user Integration to HPE GreenLake Central • Metering enhancements including show back • Simple quoting tool integration • Public, private cloud integration 	<ul style="list-style-type: none"> • HPC Customer portal integrations • Specific ISV platforms • New UoMs for HPC clusters. More granularity (core, memory, IO, bandwidth, latency and capability-based) UoM • SLA/QoS-based billing
	Initial platform	Expanded HPC reach	Exascale platform
Platform releases	<ul style="list-style-type: none"> • HPE Apollo • Standard storage • Aruba interconnect 	Cloud modules <ul style="list-style-type: none"> • HPE Parallel file system storage, Slingshot and GPUs • HPE Ezmeral Container Platform— integrated support for Slurm and Singularity • MLOps, AI integrated 	<ul style="list-style-type: none"> • Cray compute, storage options • Autoscaling of resources with any combination of standard and premium offerings • Flexible storage layers (HPE Ezmeral DF, HPE Parallel filesystem, Intel DAOS)
Key use cases	<ul style="list-style-type: none"> • Ansys/CAE • Others opportunistically 	<ul style="list-style-type: none"> • Visualization use cases; workflow • Additional workload focus EDA, CFD, FSI, Biomedical, molecular dynamics 	<ul style="list-style-type: none"> • Additional workload focus Seismic, Weather forecast, high content screening, FSI risk management

Thank you

César Gómez – cesar.gomez@hpe.com