Hewlett Packard Enterprise

Enabling Efficient and Effective Complex HPC and AI Workflows

Larry Kaplan
HPE Senior Distinguished Technologist
Chief Software Architect HPC and AI Solutions

MulticoreWorld XII 2025

Who is Larry Kaplan

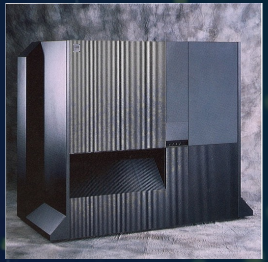

- Not the video game designer!

- Long-time HPC systems and software specialist – 30+ years in HPC
  - Started at BBN Advanced Computers working on MPP systems in 1988
    - Focus on virtual and physical distributed memory systems
  - Moved to Tera Computer Corp in 1991
    - OS engineer – virtual memory and low-level software for MTA
  - Tera bought Cray assets from SGI and renamed Cray Inc. in 2000
    - OS lead for Redstorm (Sandia)
    - Lead software architect for XT4 through XE6 starting in 2002 – added system management and programming environment
    - Named Chief Software Architect for Cray in 2013 working on XC and EX product line software stacks
  - HPE acquired Cray in 2019
    - Continuing as Chief Software Architect for HPC, adding AI and complex workflows

- Hobbies: waterskiing, snow skiing, SCUBA, science fiction

HPC and AI Supercomputers Built on HPE Expertise

Networking Slingshot

HPE leverages its vast expertise to produce open end-to-end solutions across all system components

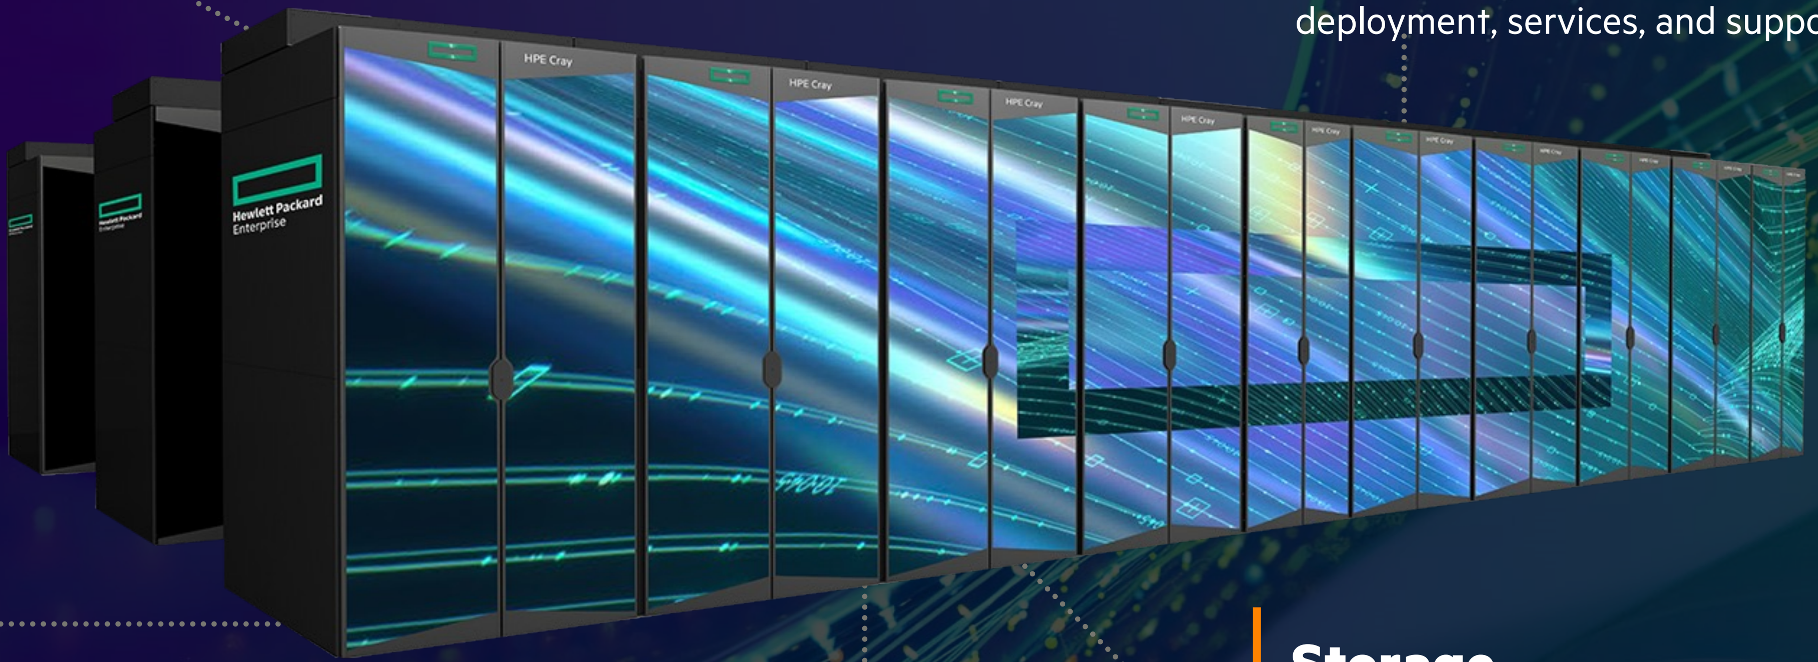

End-to-end HPC & AI Solutions

Racks and Facility Integration
World-Class Manufacturing
World-wide capability-class deployment, services, and support

HPC & AI Software

User Developer Environments
System Management, User Services

Sustainable Solutions

Cooling
Power Management

Storage

ClusterStor, DAOS

HPE Delivered Unprecedented Large-Scale HPC and AI

Helping government organizations tackle the grand challenges of humankind

37,632

GPUs

63,744

GPUs

44,544

APUs

100% liquid-cooled HPE Cray EX supercomputer

High performance GPU accelerated blades

HPE Slingshot exascale interconnect

Cray ClusterStor file systems

Enabling Large-Scaling AI Workloads Around the Globe

Los Alamos NATIONAL LABORATORY EST. 1943

10 EFLOPS

Reduced-precision AI Performance with NVIDIA GH200 superchips

CSCS Centro Svizzero di Calcolo Scientifico Swiss National Supercomputing Centre

20 EFLOPS

Reduced-precision AI Performance with NVIDIA GH200 superchips

University of BRISTOL

21 EFLOPS

Reduced-precision AI Performance with NVIDIA GH200 superchips

Complex Federated Workflows

HPC-Analysis Amalgamated Workflow

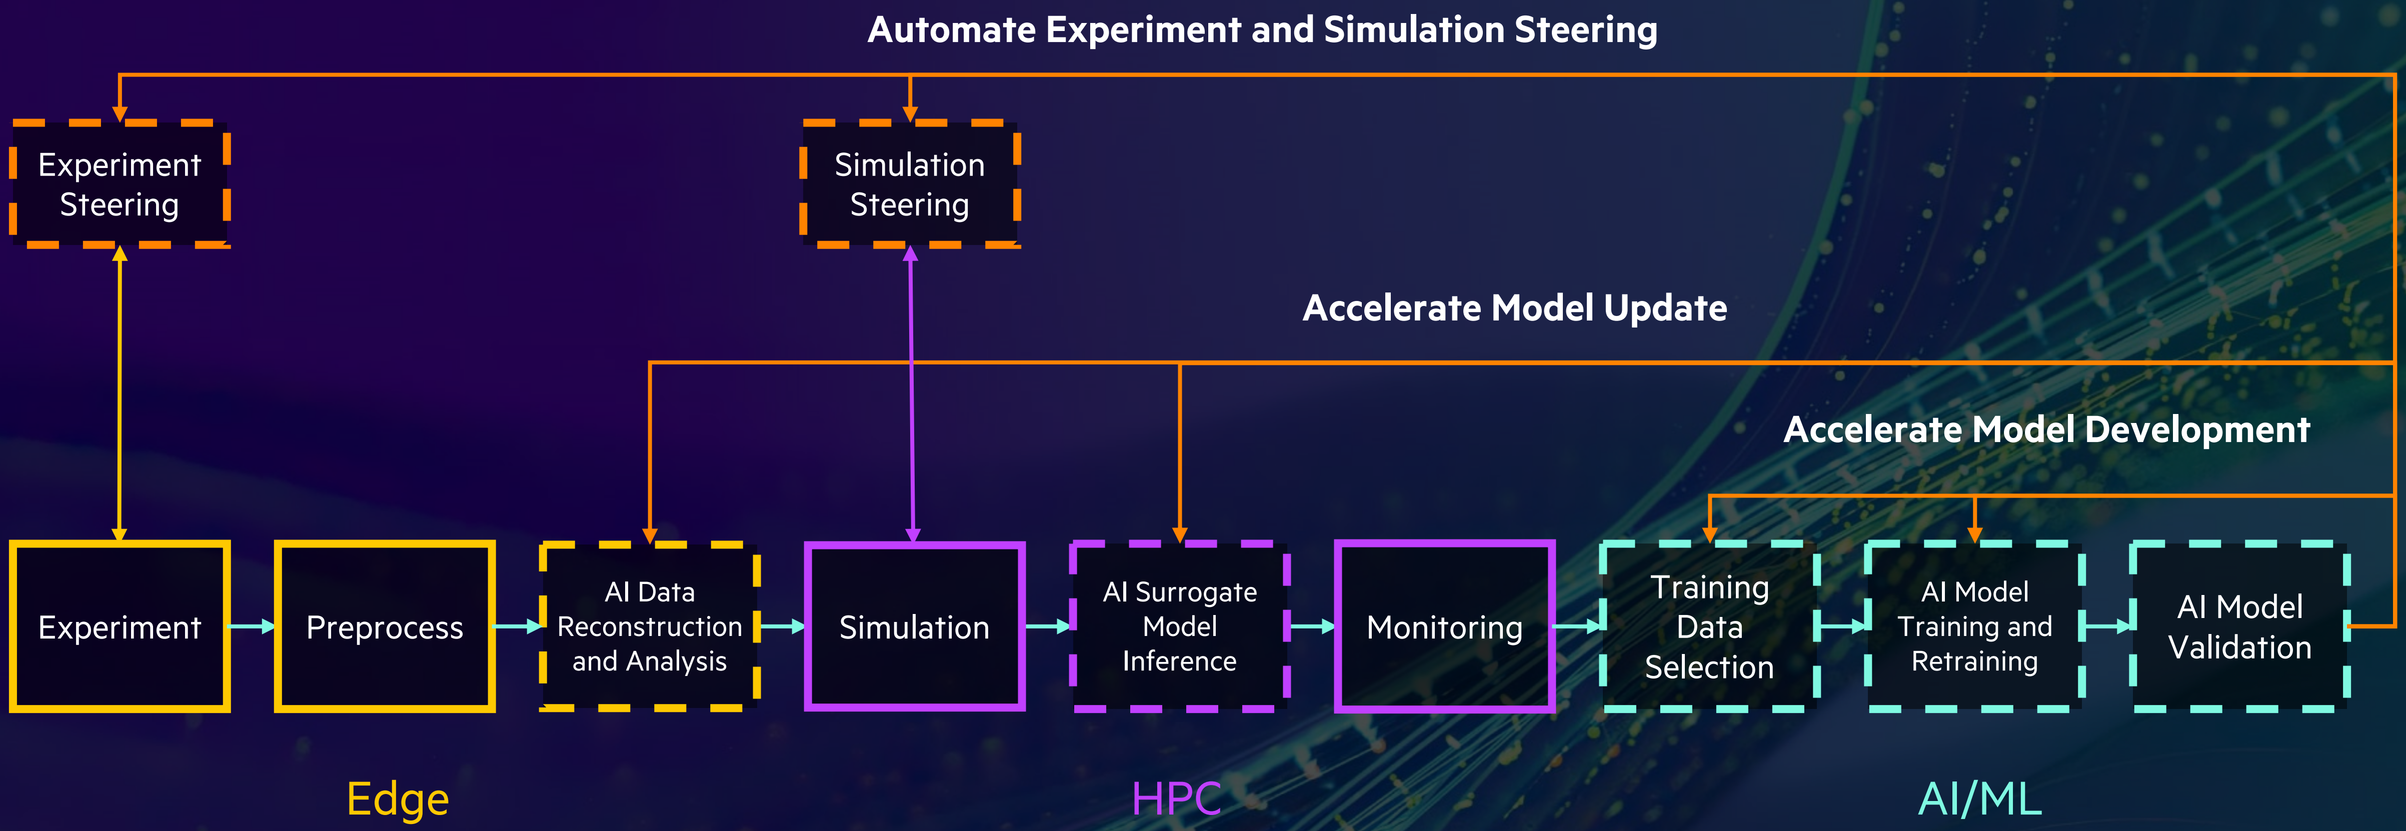

HPC AI Software: Efficiency and Composability

Infrastructure

- Desire for **flexible access**
- Use of as-a-service models, especially platform (**PaaS**) and software (**SaaS**), potential for Serverless computing (**FaaS**)
- **Holistic** Power and Energy management with TCO benefits

Runtime

- Common community use of **containers** for packaging and delivery
- Desire for improved **container ecosystem support**
  - Build tools, registries
  - Choice in orchestration (batch/Kubernetes)

Workflow

- Standardized **APIs** for composition and deployment across federated diverse nodes and systems
- Ability to **schedule** disparate resources
  - Processors/nodes, networks (overlays, QoS), storage
- Supporting **services** of interest
  - Databases, metadata handling, workflow tracking

Customer Concerns for Power and Energy

Optimized Job Performance under resource constraints

- Run hardware over-provisioned systems with better overall system performance
- Need: balance between available power/cooling and workload performance

Minimize Energy Consumption

- Reduce OPEX in the face of increased power needs of new technologies and increased energy prices
- Need: reduce energy consumption of workloads according to time–to-solution vs energy–to-solution tradeoffs

Power and Energy Management

Data Center Sustainability

- Governments (US, EU) are developing mandates for our customers to address sustainability
- Need: bring down current energy usage and carbon footprint, optimize system operation according to data center TCO (balance facility efficiency with system operation)

Maximize Resource Utilization

- Optimize power and cooling needs to support sustainable HPC efforts
- Need: optimized use of available resources (including heat reuse), minimize stranded capabilities (power, cooling) in datacenter and HPC systems

What is HPE Considering?

Flexibility

Enable Diverse HPC/AI Workflows
with easily deployed services and uniform APIs

Containerization
enables consistent, portable applications and environments across varied infrastructures, virtualized where appropriate

Multiple Linux Distributions
including Red Hat Enterprise Linux and other popular distributions

Sustainability

System-wide Power Management
reduces energy/power CapEX and OpEX by steering power to jobs/components on demand

Application-Level Optimization
trades off energy savings for minor performance impacts under user control (with no burden)

HPE Software
to revolutionize user experience

Efficiency

Complex Workflow Orchestration
facilitates execution of diverse workloads on HPC infrastructure

Storage and Transfer Operations
address issues of data movement and staging

Low Noise Mode
focuses processing resources on running workloads

Virtualization

Proposed for medium-term

- Integration with virtualization tools for virtualized user services
  - Virtualized compute and service nodes
  - Provides environment isolation

Under investigation

- Towards confidential computing
- Virtualized user services for SaaS
- FaaS for fine-grained decomposition

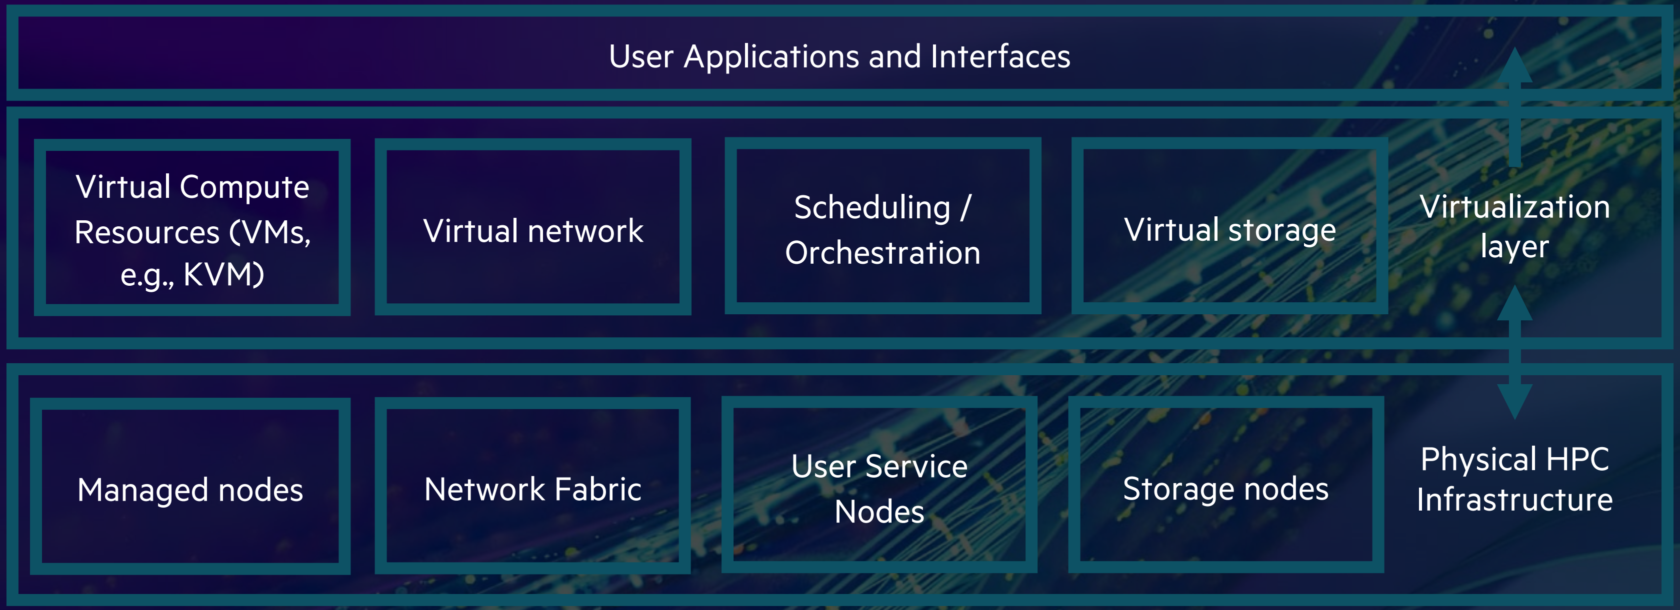

Workflow Services

Proposed medium-term

- Workflow state tracking
  - Private databases (e.g., for CMF, visualization)
  - Private event bus (pub/sub)
  - Tracking for transfers, network, and storage in addition to compute
  - Provide ease of deployment and use of community solutions

- Service resiliency

Under investigation

- Data management
- Extended workflow tracking and management
  - Includes events across sites and facilities (e.g., data readiness at experiment/instrument)
- Additional components
  - Example: workflow composition tools – environments, debugging, profiling
- Integrated QoS – across processing, networks, storage

Runtimes for AI and HPC

Dragon is a composable distributed runtime that enables users to create sophisticated, scalable, resilient, and high-performance AI/HPC applications, workflows, and services through standard Python interfaces

Distributed Python and interactive Jupyter	Resilient training and inference	Easy-to-use and HPC-optimized data exchange between applications	Observability for AI/HPC workflows
Multiprocessing	AI	Data	Telemetry

Dragon Core

- 2 – 100X faster data processing than Ray
- Scalable to over 1000 nodes
- Multi-system features offer a hybrid experience, spanning from laptop to supercomputers
- Open-source or HPE-optimized packages
- Well-documented with numerous cookbook examples and easy setup

developer.hpe.com/platform/dragonhpc/home

github.com/DragonHPC/dragon

Common Federation Framework: Workflow Deployment SDK

Enables federated hybrid workflows on data from Edge to Extreme Scale to Cloud

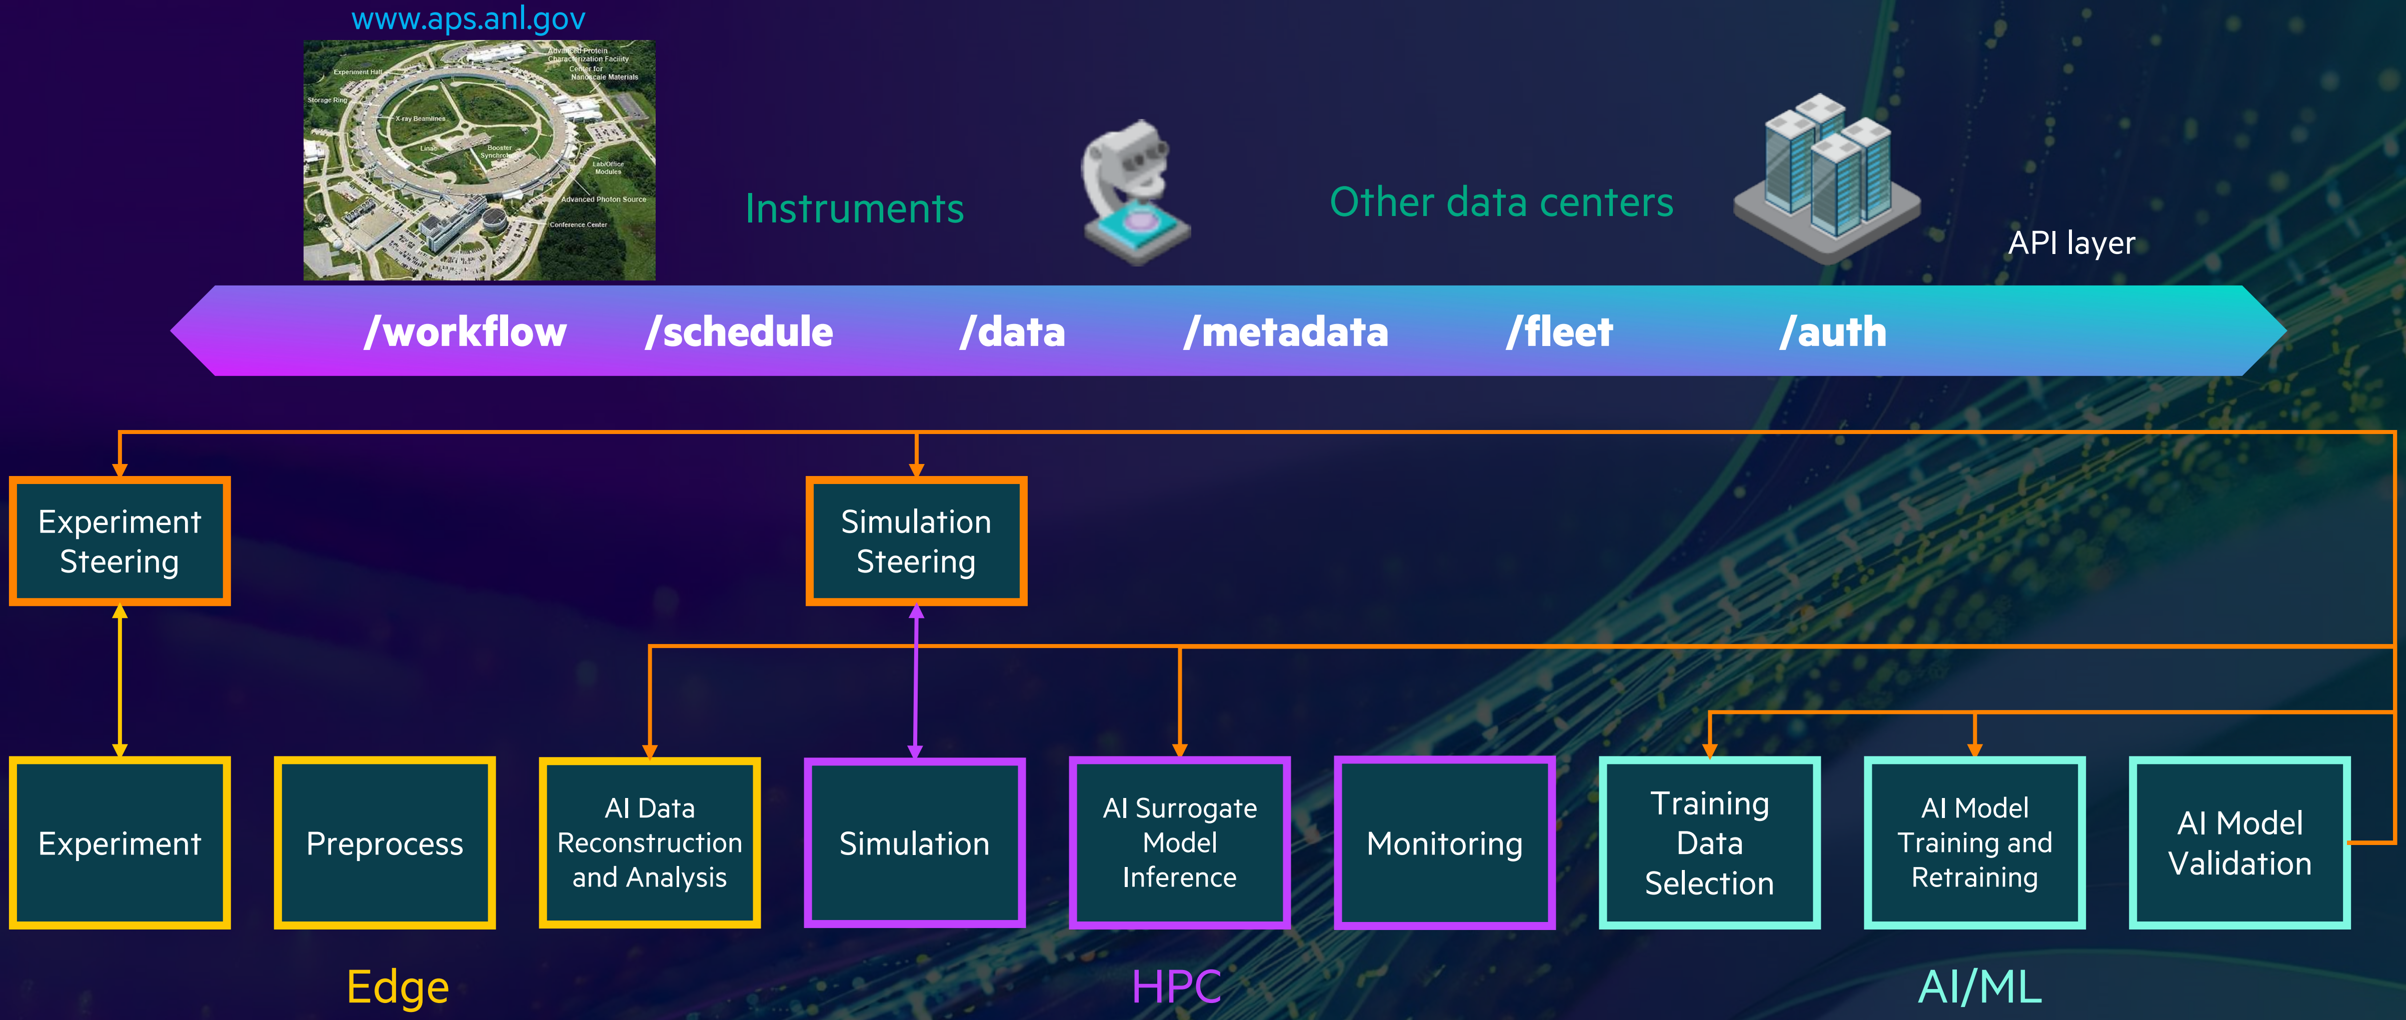

Advanced Power and Cooling Management

More resilient systems and better datacenter economy

- **Scope of power and cooling**: nodes (CPU and GPU), chassis, rack AC, bulk DC, liquid-cooling infrastructure
- **Power monitoring and analysis**: Read and aggregate power and energy at available measurement points for analysis and planning, both in and out of band
- **Topology-aware node power on/off and steering**: Step-by-step topology and protocol-aware power management enables controlled operation
- **Power management of jobs**: Constrain power consumption and apply power per user, per group, and per job (interface with Slurm and Altair PBS Professional) and account for all consumed energy
- **Power caps**: Set limits to trigger power cap based on environmental failure (power or thermal), data center power capacity, or other reasons like planned brownouts, time of day
- **AIOPs**: anomaly forecasting for facility metrics (cooling devices) – 15 minutes ahead of issues

Monitor & Manage Power | Analyze & Forecast Power | Cap Power | Manage Workloads

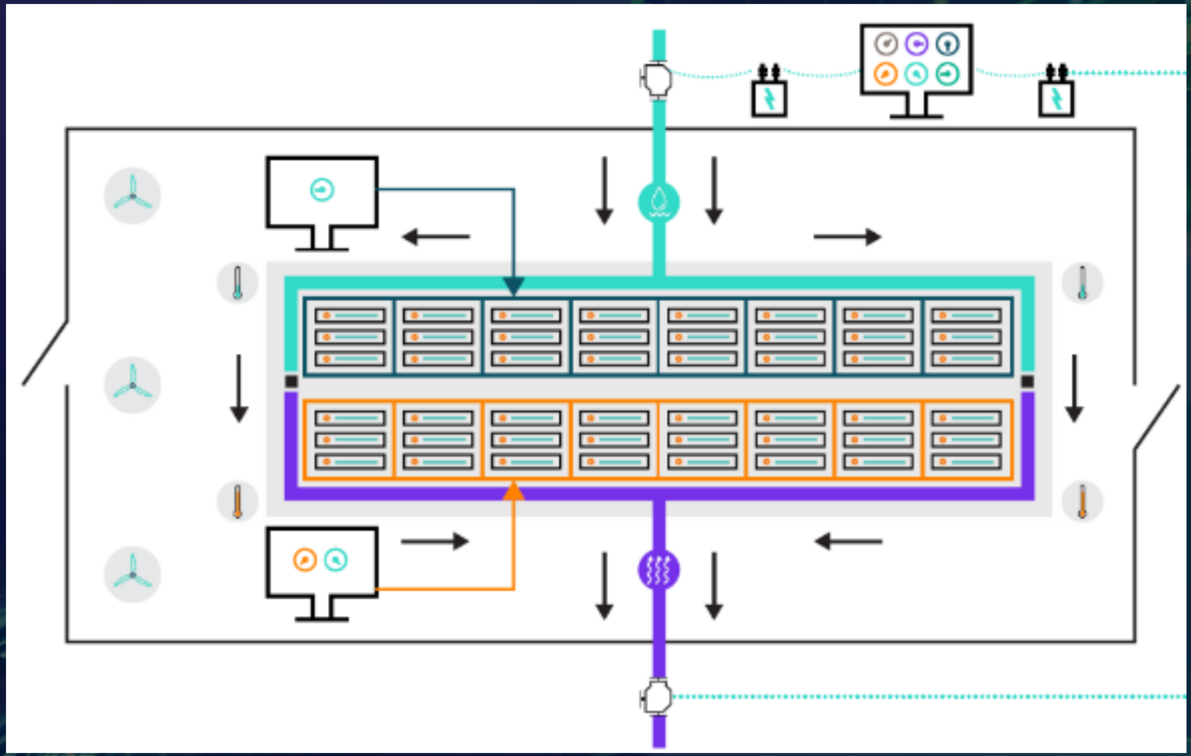

Holistic power and energy management
Enables more science per Watt*

Concept: System Administrator and/or User define optimization policy

Minimum resources ⟷ **Maximum performance**

Holistic power and energy management tools

HPE's automatic workload control delivers desired outcome

Possible constraints
Power
Cooling
OPEX
Sustainability

Monitoring
- Data
- Events
- Decisions

- Dynamically balance between available power/cooling, optimized resource usage, and workload performance
- Balance facility efficiency and system operation with minimal performance impact

* Potentially up to 50+% power and TCO savings

A TCO savings example

Estimates for 8 theoretical racks (each 200kW IT nameplate power)	No Management	Uniform Static	Strategy 1	Combine (Strategy 1&2)	Combine (Strategy 1&2)
Application Performance	100%	>90%	>99.1%	>95%	>90%
IT compute power (MW)	1.6	1.2	1.2	0.9	0.7
Facility Power procured (MW)	2.3	1.7	1.7	1.2	1.0
OPEX 5 years (Million US)	>=8.4	8.4	8.4	6.0	5.0
CAPEX savings (Million US)	0.0	**4.3**	**4.3**	**7.5**	**9.3**
OPEX savings over 5 years(Million US)	0.0	0.0	0.0	**2.4**	**3.4**
Potential annual OPEX savings (Million US)	0.0	0.0	0.0	**0.5**	**0.7**
Perf/procured Watt efficiency (relative)	**1.00**	**1.23**	**1.35**	**1.79**	**2.14**

HPE Software Innovation – Please provide feedback on areas of interest

- Complex workflows and AI support
  - OS distributions and support
  - Scheduling (batch and other)
    - State tracking and events
  - Containers/virtualization
    - Portability and isolation
  - Services
    - Data management – private databases, data intelligence, metadata and provenance
    - Authentication/Authorization Person/Non-Person
  - Unified API for workflows

- Power and energy management
  - Use cases
    - Efficient operation, free cooling, excess heat usage, etc.
  - Lower TCO (both OpEX and CapEX)

Thank you