

Beyond the Horizon: Quantum, AI, and the tech shaping the next decade

Andrew Wheeler, HPE Fellow, SVP and Director, HPE Labs

MulticoreWorldXIII, February 2026

The Shape of the World in 2030

1



**Agentic
Enterprise**

2



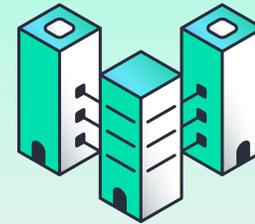
**Beyond
the Grid**

3



**Proof
Before Trust**

4



**Quantum
Utility**

5



**AI
After LLMs**



The Shape of the World in 2030

1



**Agentic
Enterprise**

2



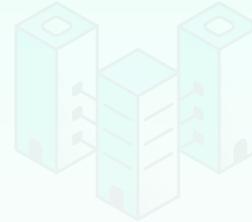
Beyond
the Grid

3



Proof
Before Trust

4



Quantum
Utility

5



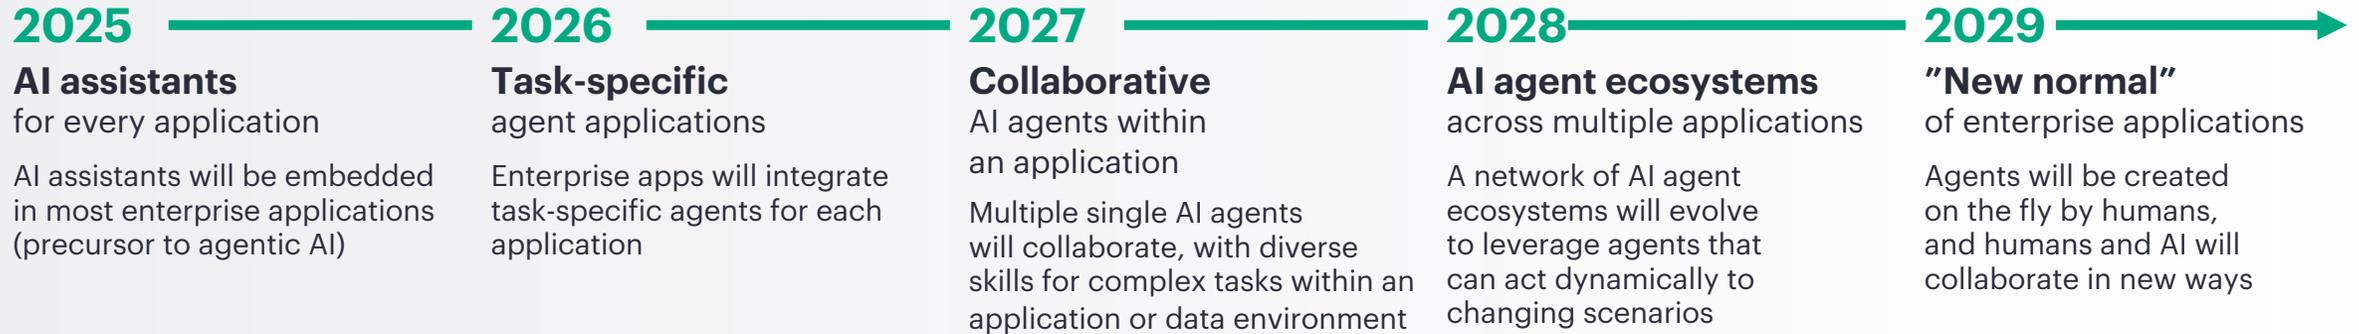
AI
After LLMs



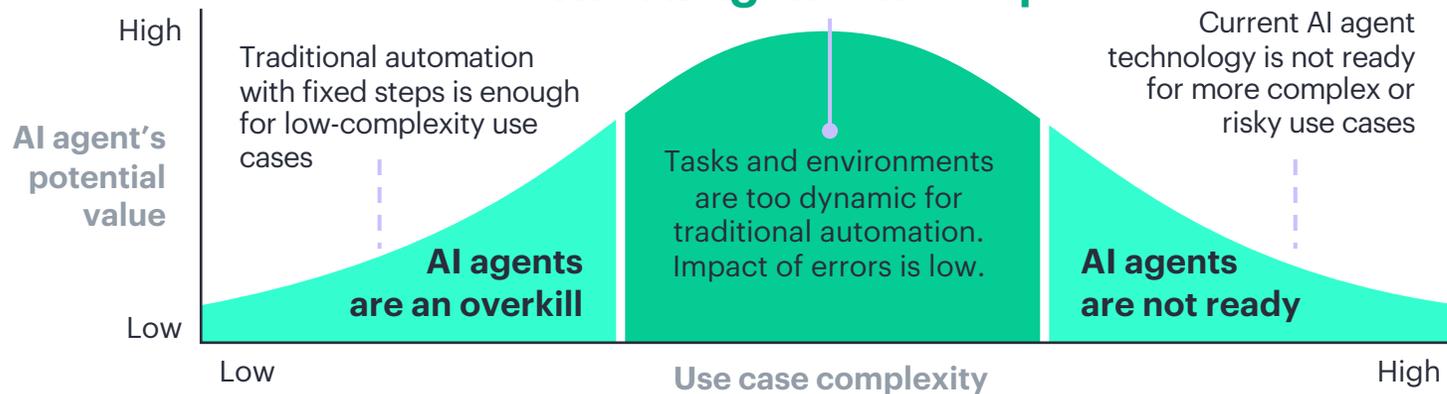
Agentic AI and the hype cycle: Signal vs. Noise

The peak of inflated expectations

The future of agentic AI in enterprise applications



The AI agent sweet spot

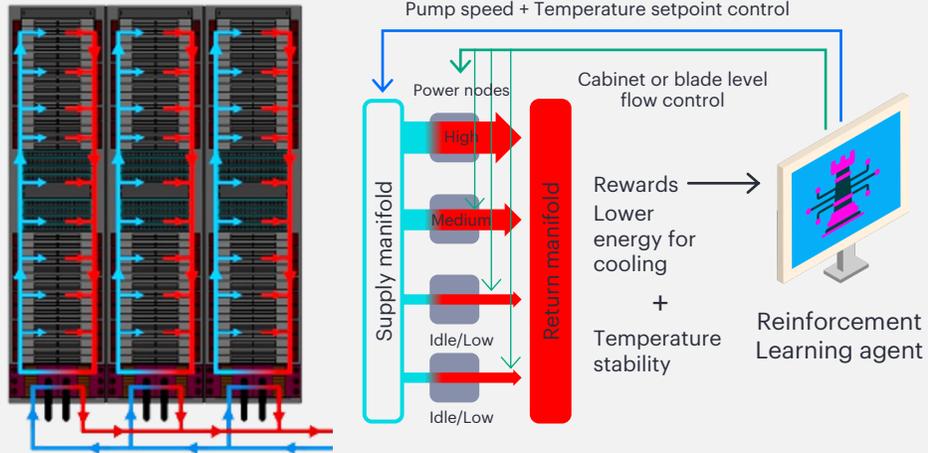


By 2035, agentic AI will drive approximately **30%** of enterprise application software revenue, surpassing **\$450B**, up from 2% in 2025

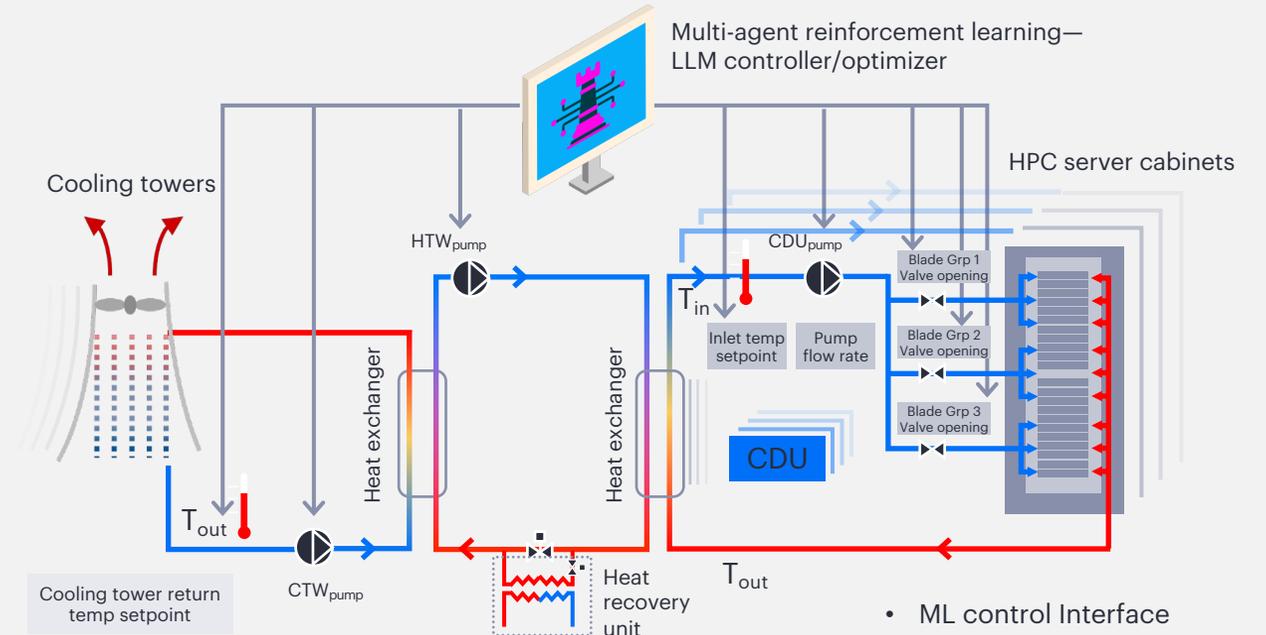
Gartner predicts that over **40%** of agentic AI projects will be canceled by the end of 2027

Precision liquid cooling for the AI era

Scalable and granular control of the Cooling Distribution Unit (CDU) for IT cabinets



Scalable cooling control of different components of a data center



- High-fidelity cooling model of data center Digital Twin
- Model scales to HPC cluster

- End-to-end model spanning cooling Tower to IT cabinets and heat recovery unit

- ML control Interface
- Multi-agent Reinforcement Learning real-time control
- Optimizes energy and temperature stability

LLM-based Agentic control plane that can reason, monitor, and explain its actions in real-time

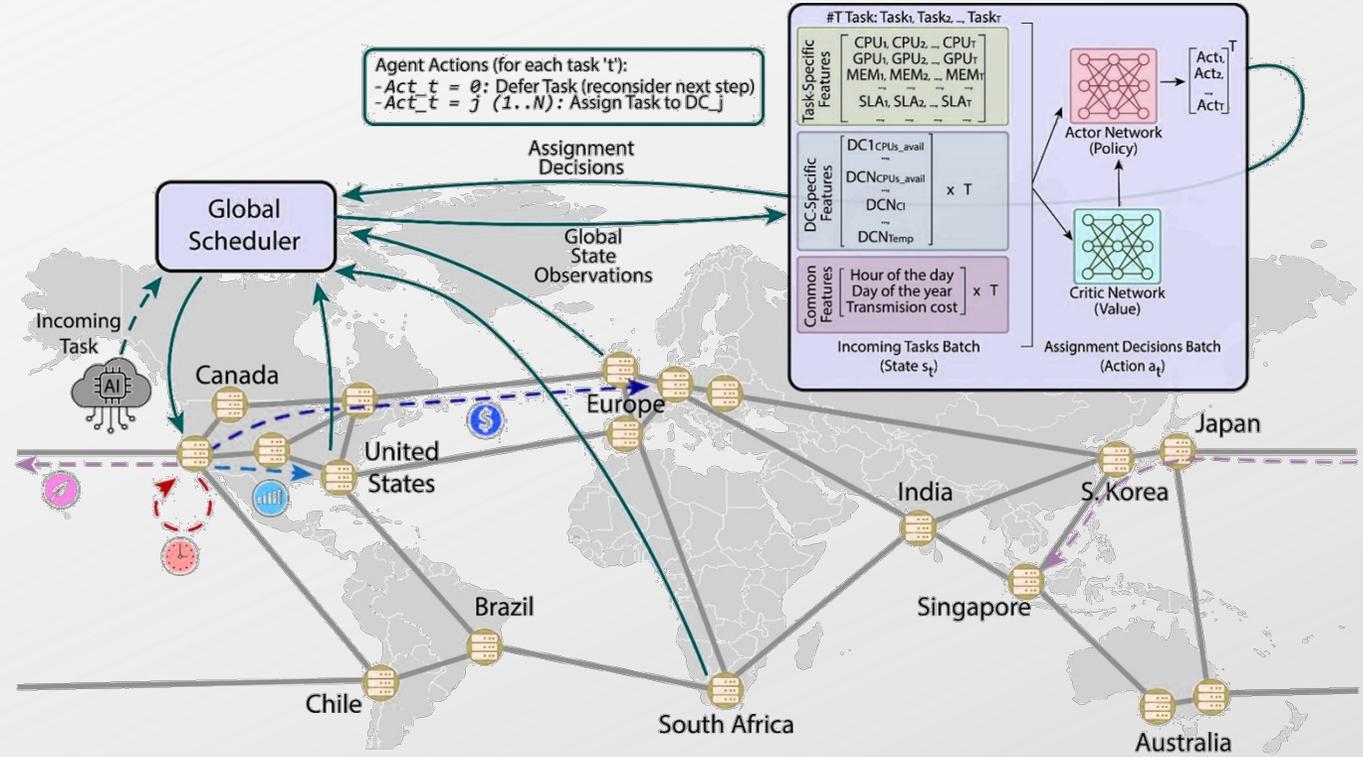
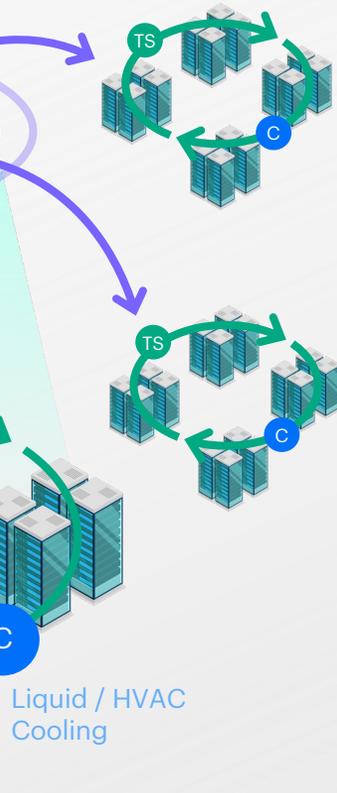
Global cloud orchestrator

Data center cluster level (global) control

Geographic shift of workload

Time shift of workload

Data center level (local) control



Agentic AI pipeline



**GreenLake
Intelligence**



**HPE Agentic
AI Mesh**



**Active
Learning AI
Agents**



**Agentic Liquid
Cooling**



**Data Center
Cloud Control**

HPE
Evolutionary

HPE Labs
Disruptive

Learn More

NeurIPS
Liquid Cooling



NeurIPS
Data Center Cloud



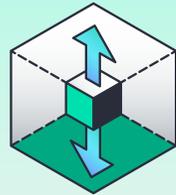
The Shape of the World in 2030

1



Agentic
Enterprise

2



**Beyond
the Grid**

3



Proof
Before Trust

4



Quantum
Utility

5

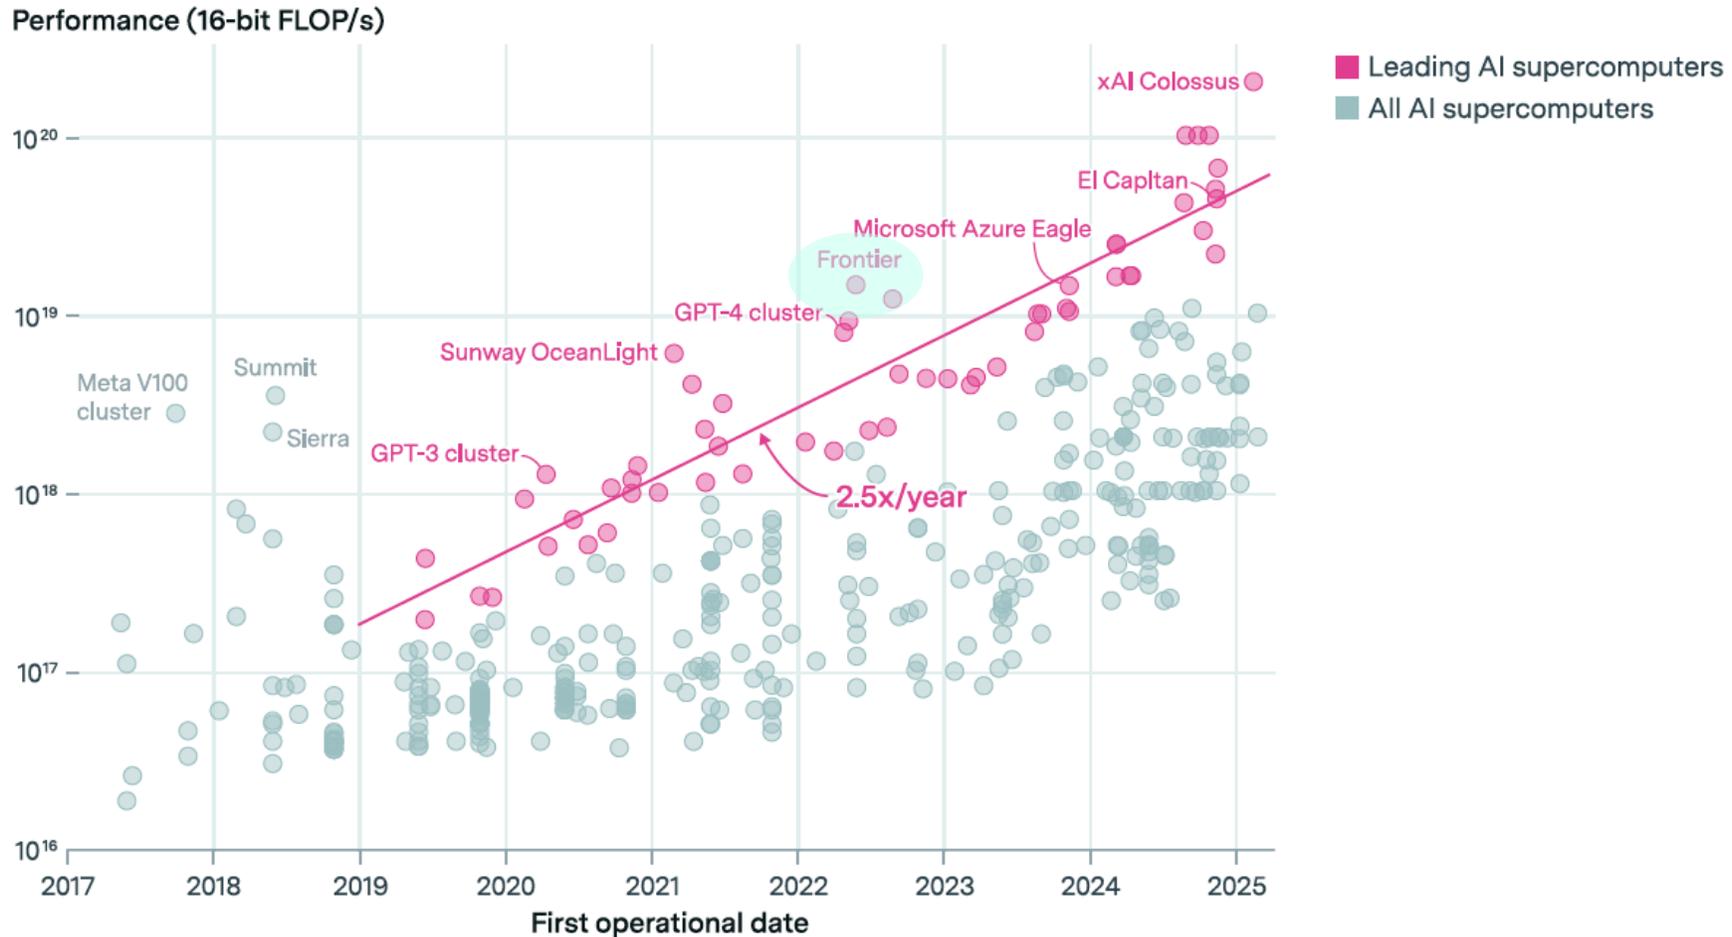


AI
After LLMs



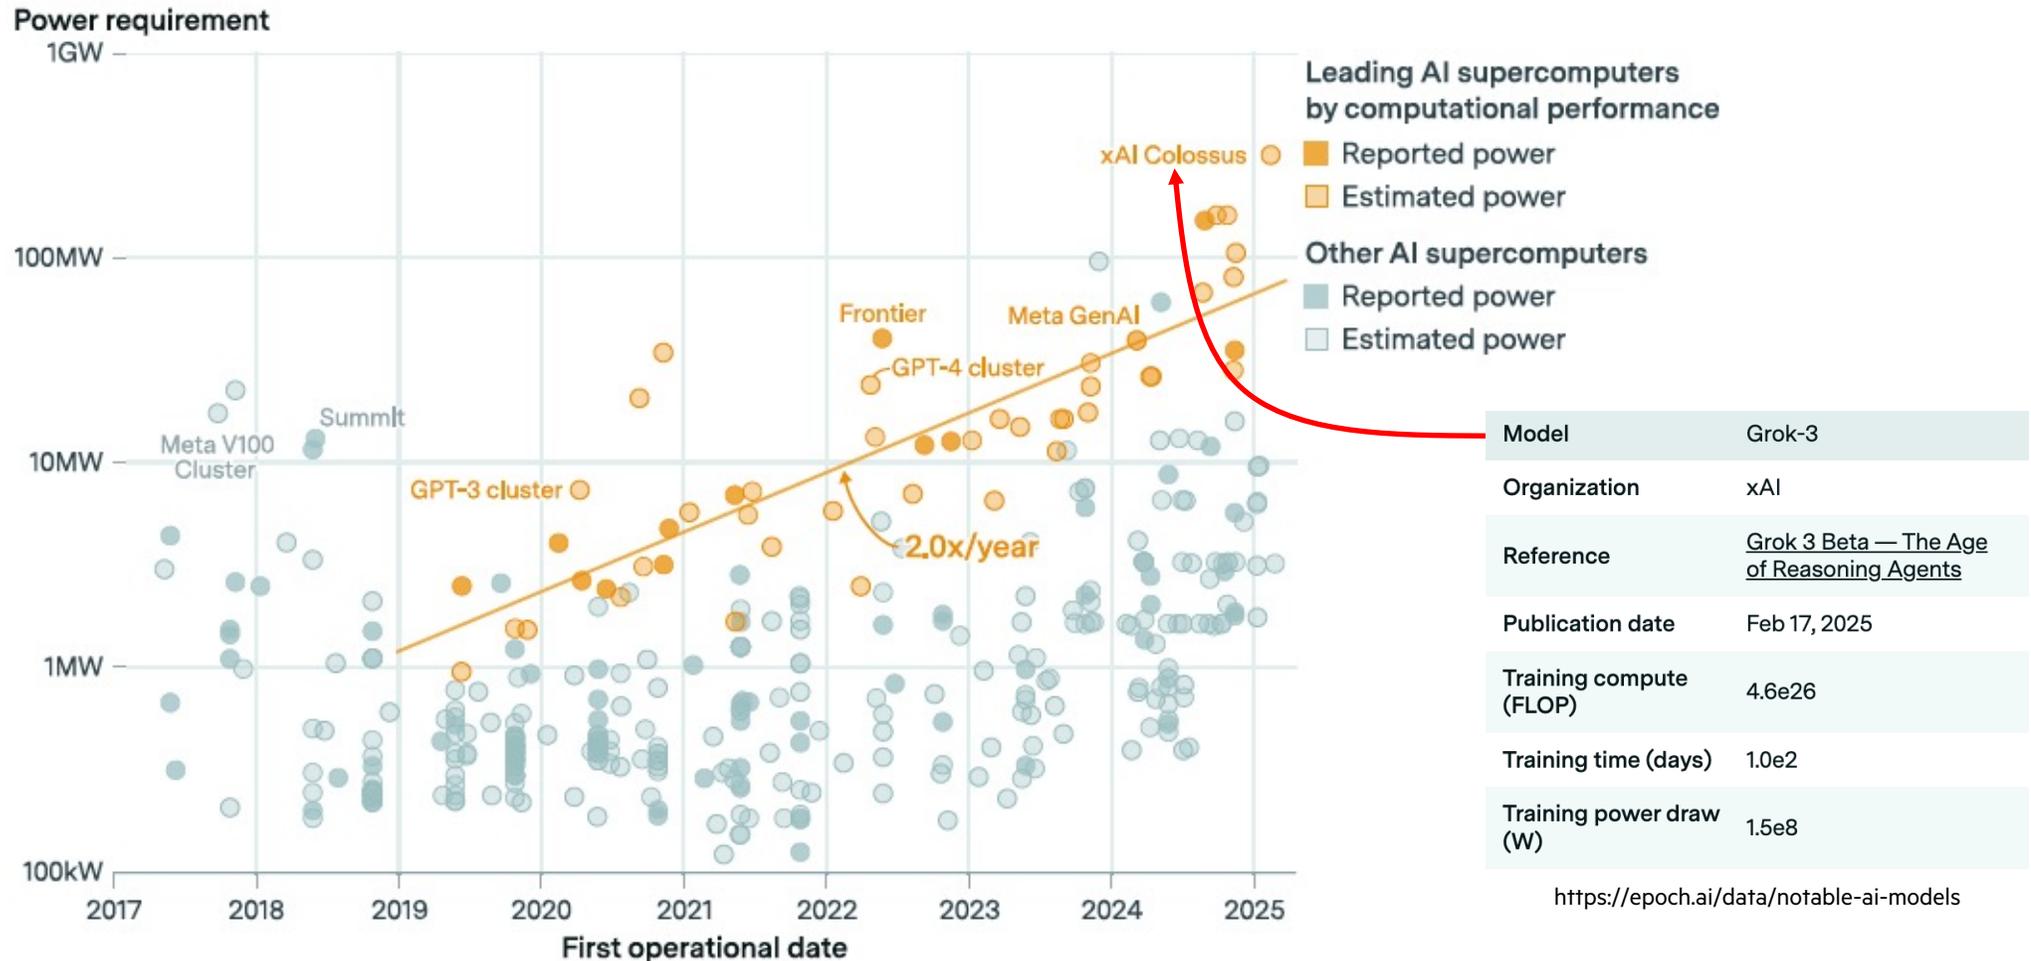
Post-Exascale Trends - Performance

Computational performance of AI supercomputers



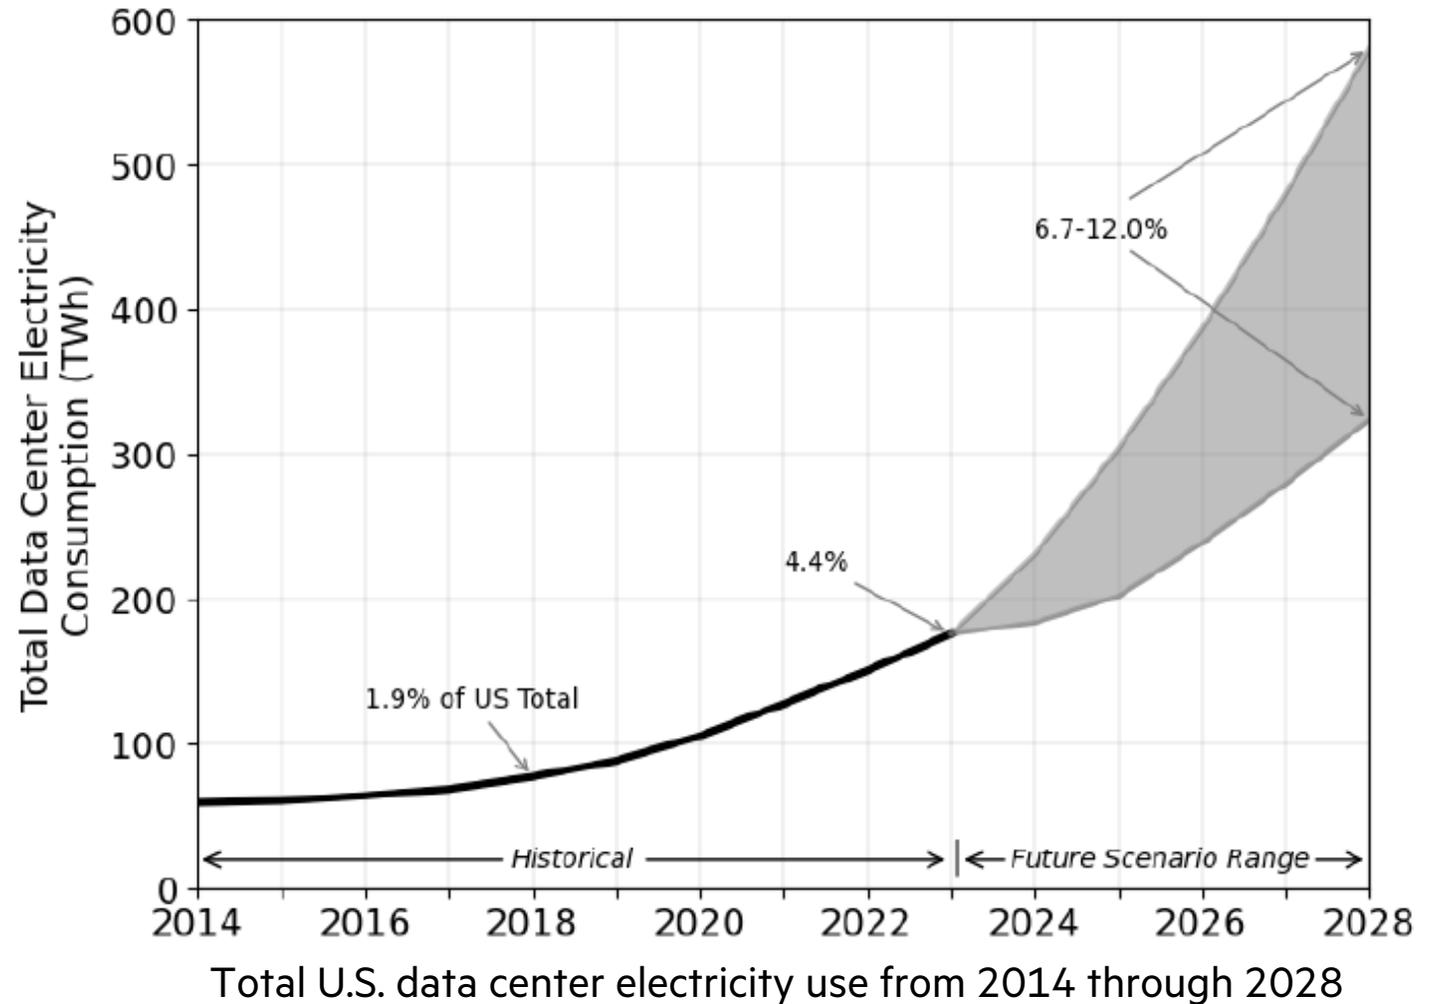
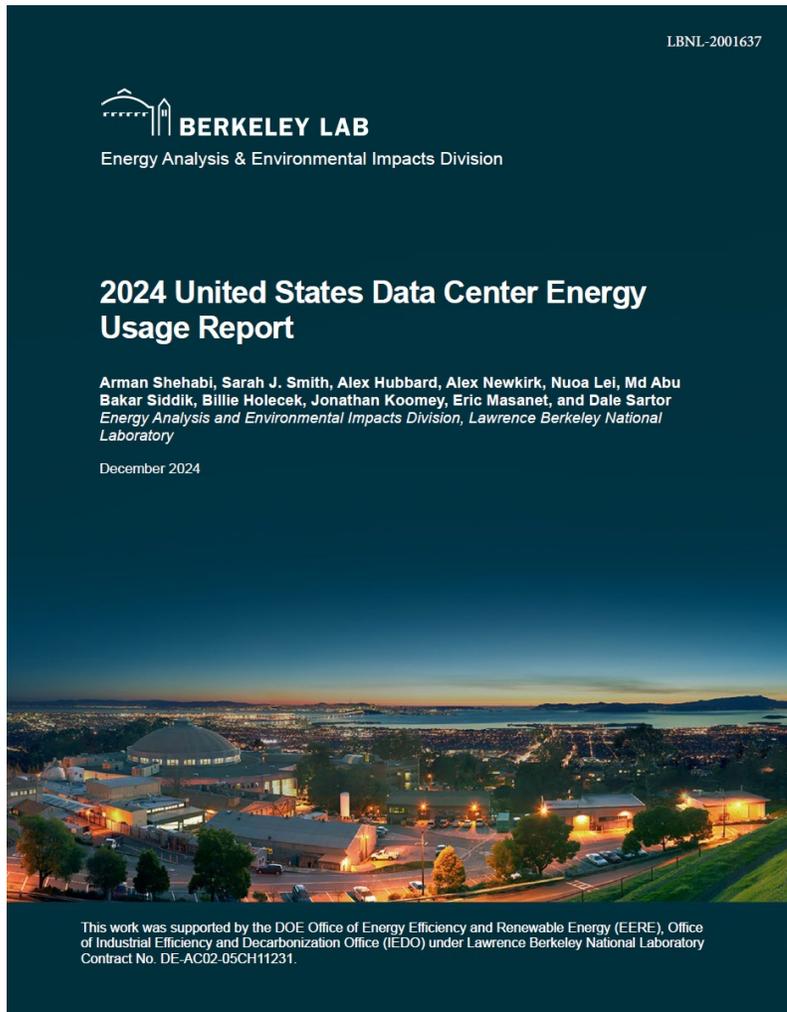
Post Exascale Trends - Power

Power requirement of AI supercomputers



Pilz, K., Sanders, J., Rahman, R., and Lennert H. Trends in AI Supercomputers, 2025

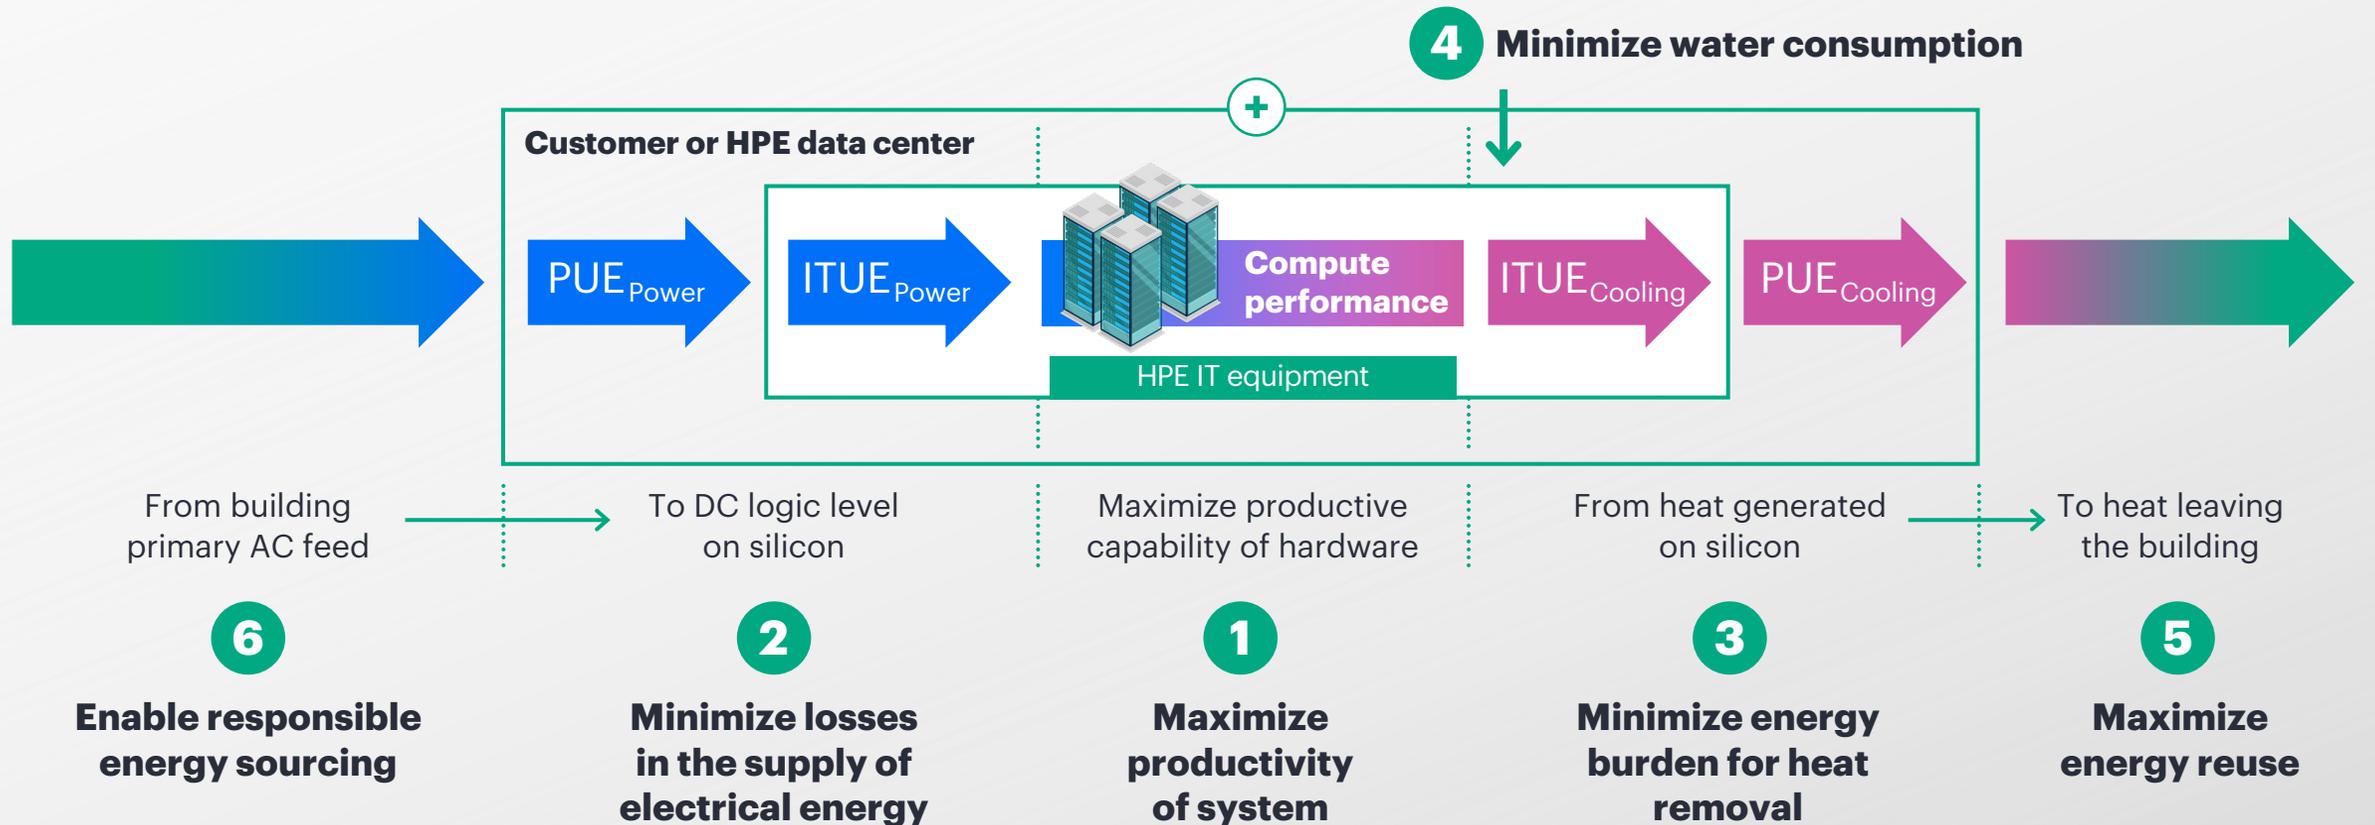
United States Data Center Energy Usage Projections



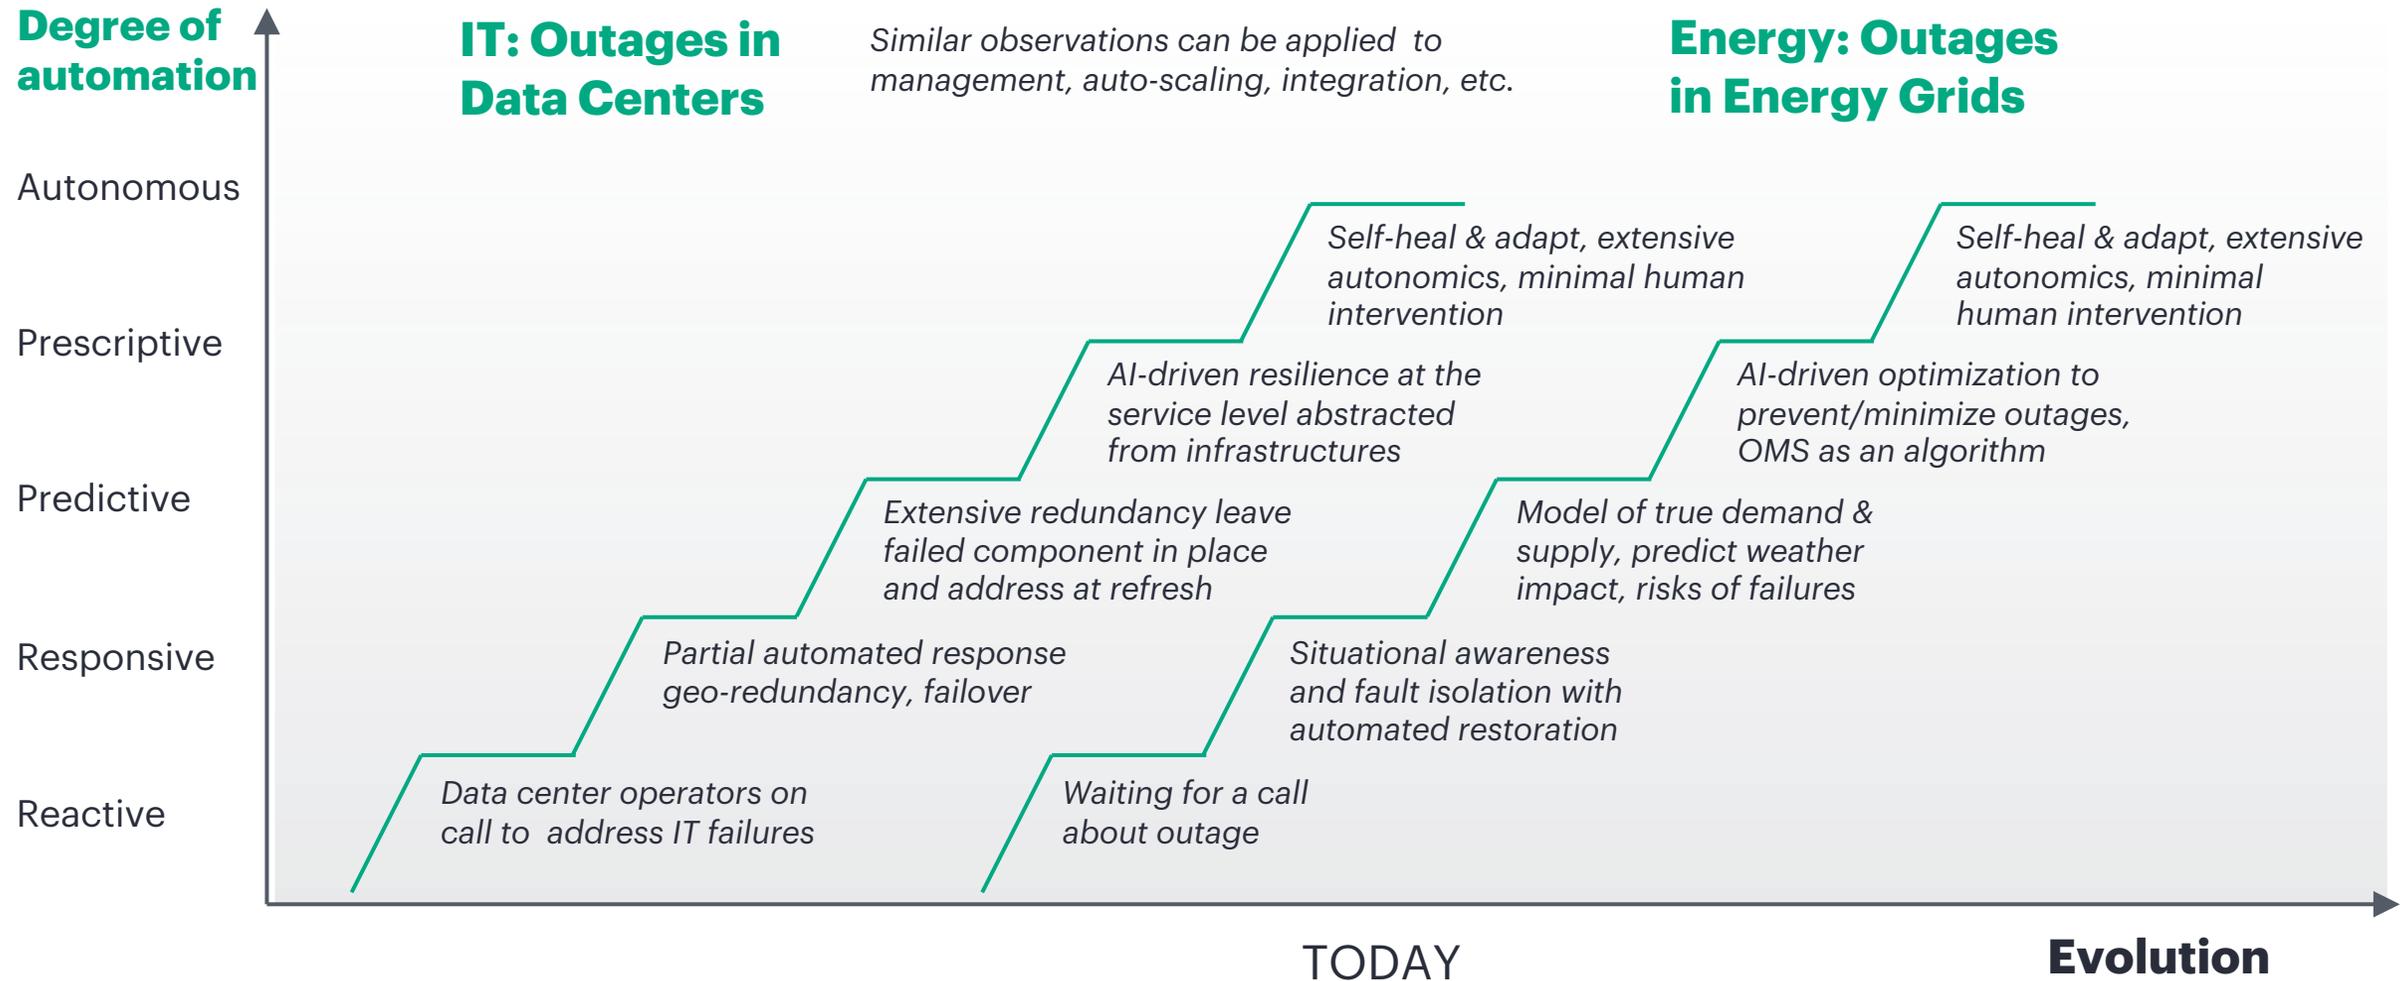
Shehabi, A., Smith, S.J., Hubbard, A., Newkirk, A., Lei, N., Siddik, M.A.B., Holecek, B., Koomey, J., Masanet, E., Sartor, D. 2024. 2024 United States Data Center Energy Usage Report. Lawrence Berkeley National Laboratory, Berkeley, California. LBNL-2001637

The Problem (and Solution space)

Data center energy flow

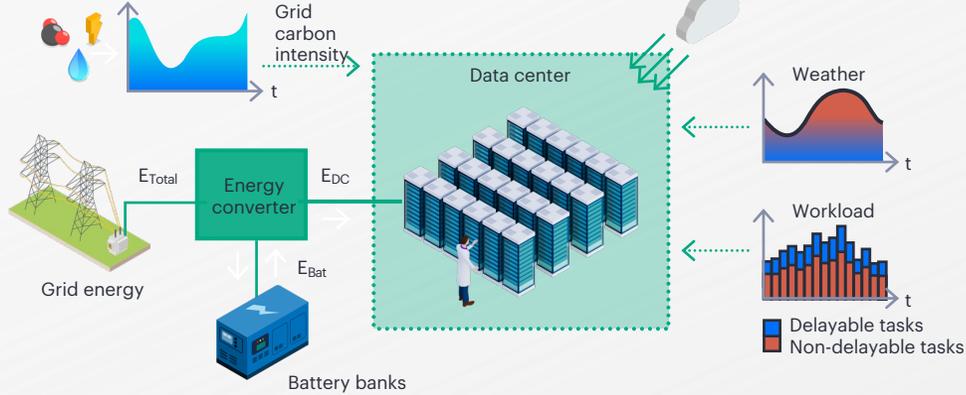


Evolving outage management: IT and energy grids



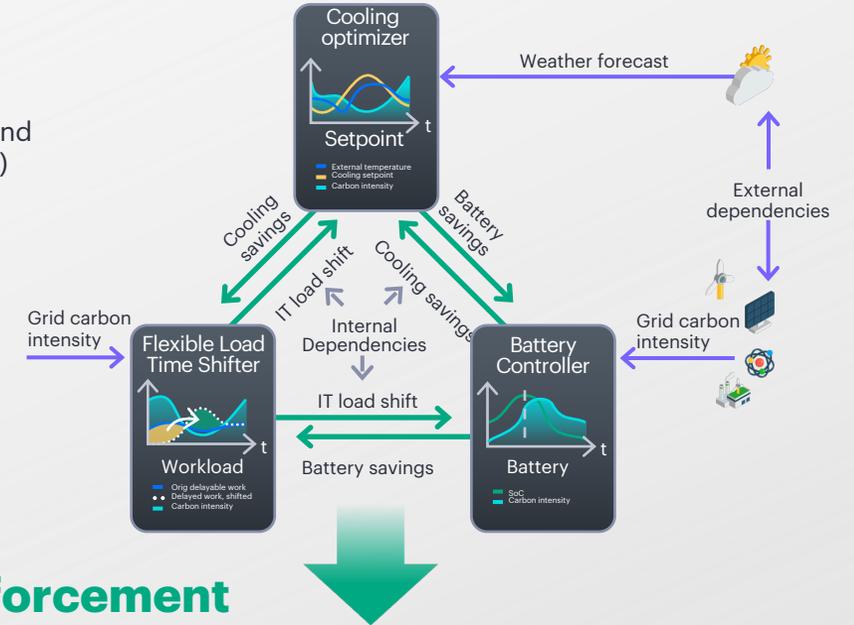
Holistic data center intelligence

Physical system

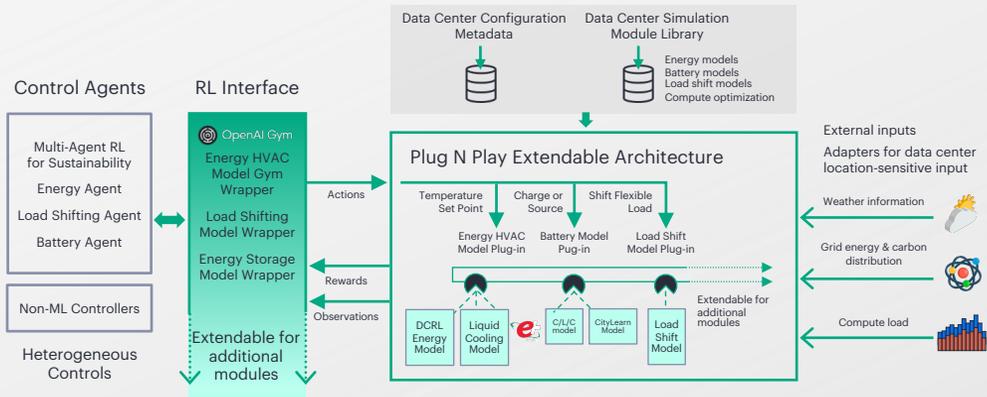


Digital system

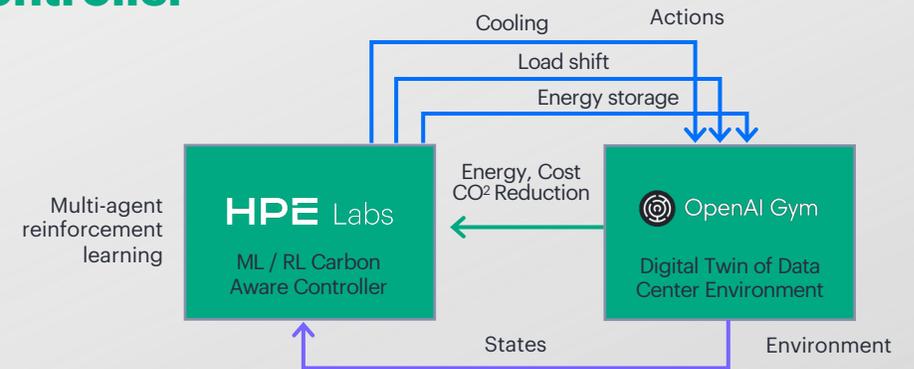
Resolving internal and external dependencies for the cooling, and scheduling (and energy storage)



SustainDC RL Environment Framework



Multi-agent reinforcement learning controller



Beyond the Grid pipeline



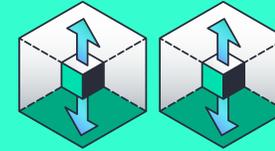
Direct liquid cooling



Sustainability Insight Center



Holistic Power Management



Digital Twins



Next-Gen Grid

HPE
Evolutionary

HPE Labs
Disruptive

Convergence of Energy and IT Sectors

Learn More



The Shape of the World in 2030

1



Agentic
Enterprise

2



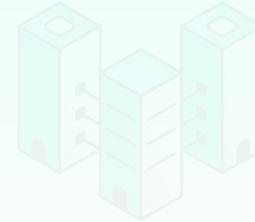
Beyond
the Grid

3



**Proof
Before Trust**

4



Quantum
Utility

5

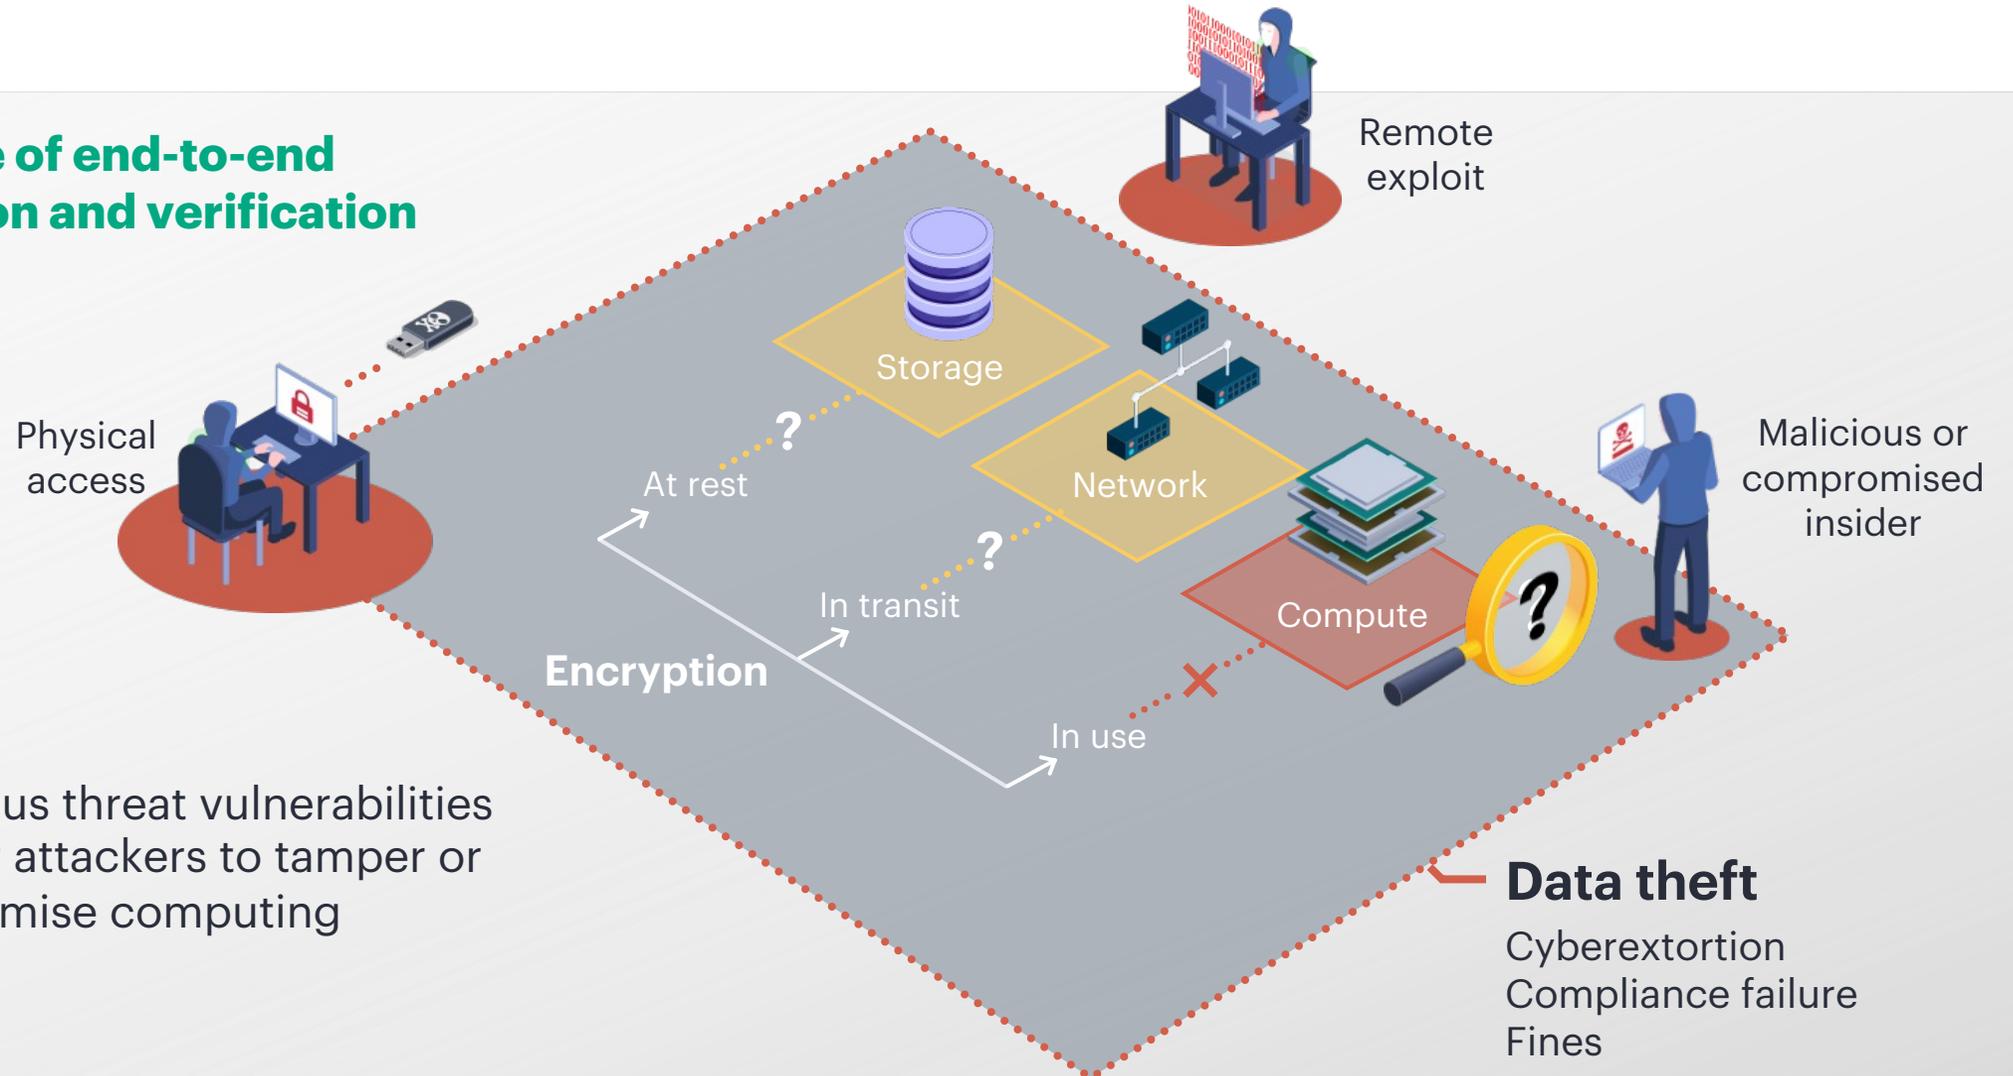


AI
After LLMs



The Problem

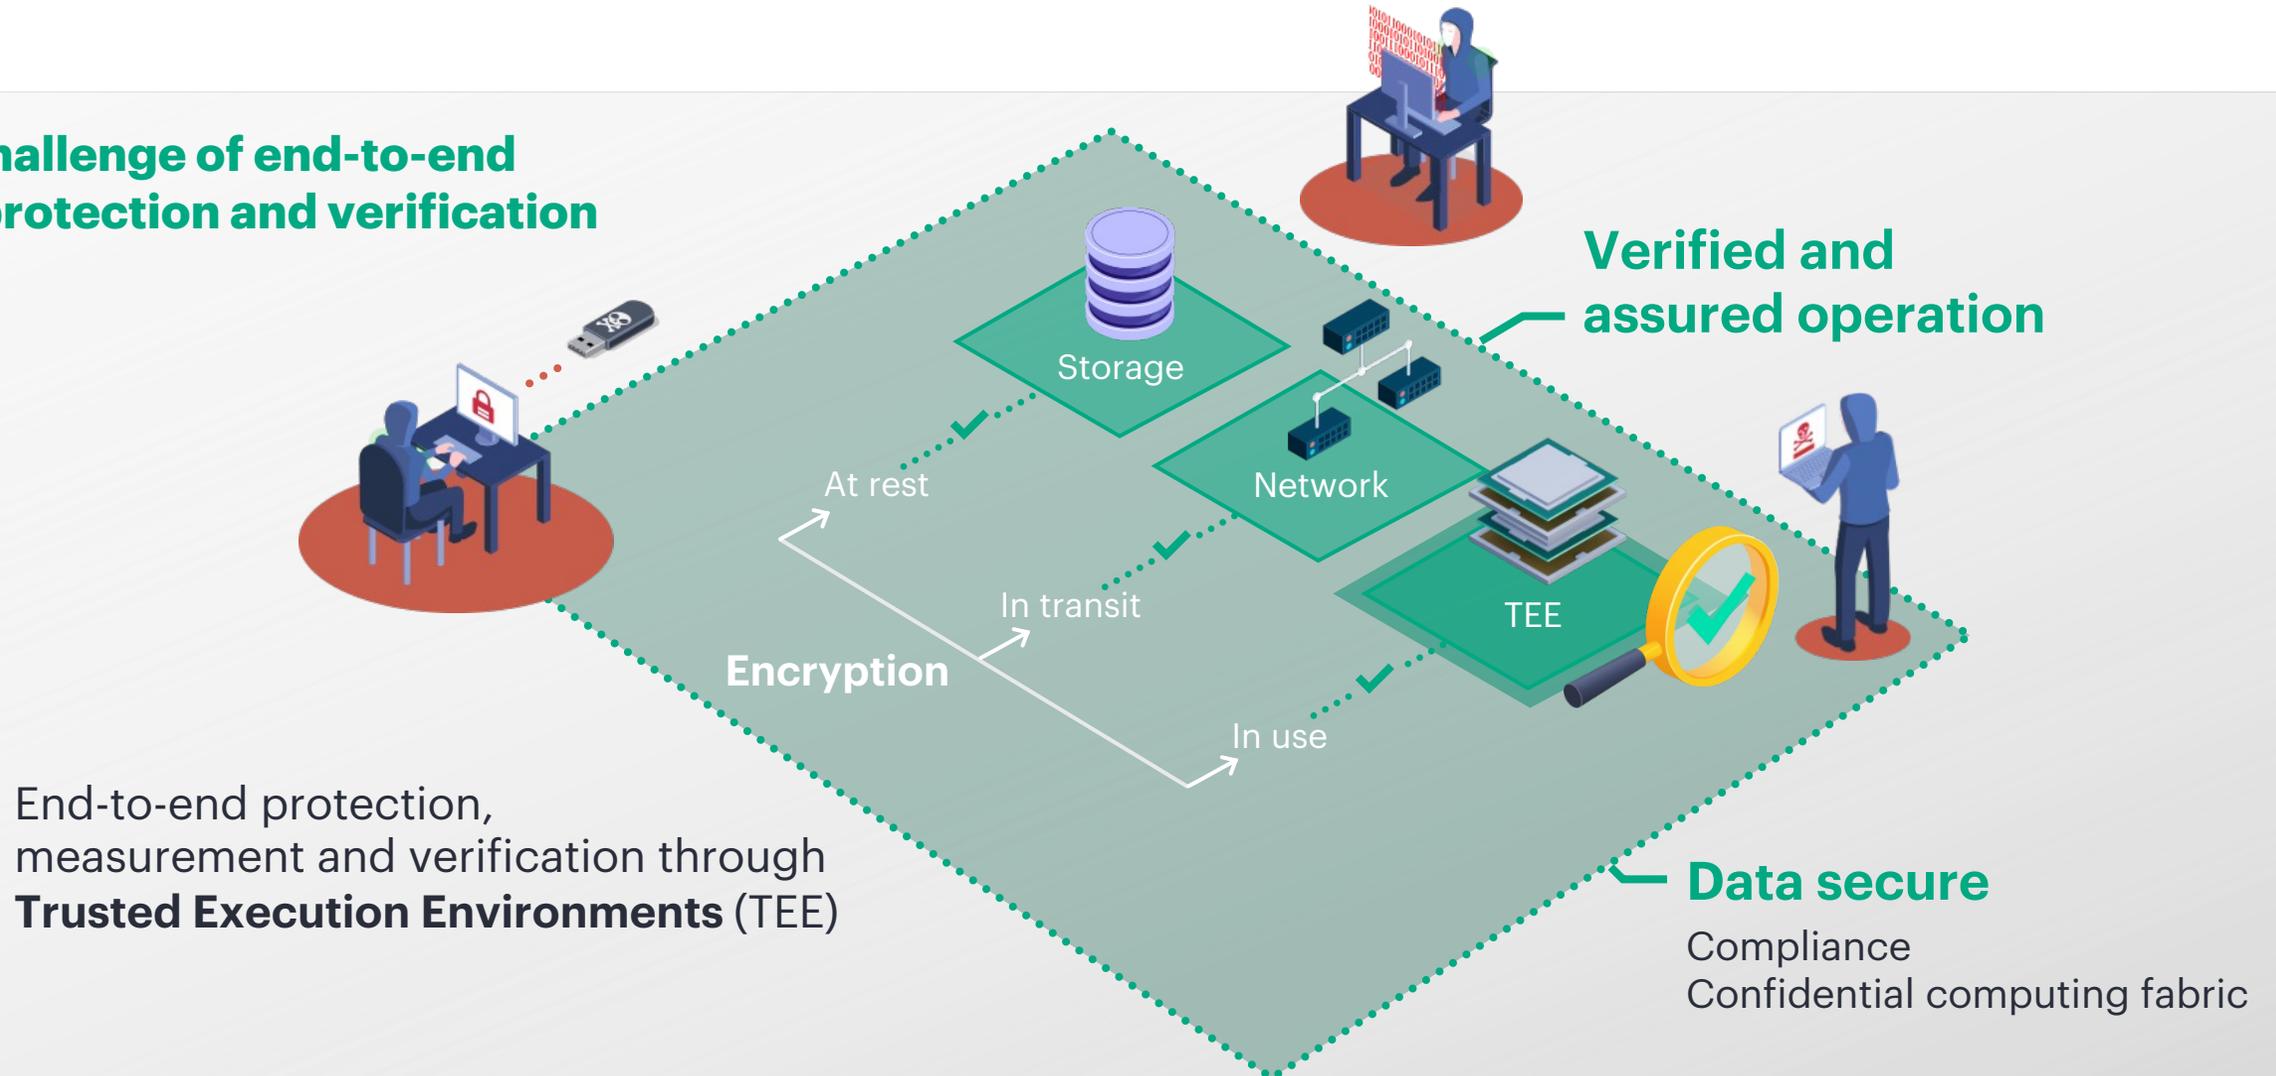
The challenge of end-to-end data protection and verification



Numerous threat vulnerabilities exist for attackers to tamper or compromise computing

What we're solving with Confidential Computing

The challenge of end-to-end data protection and verification



The Shape of the World in 2030

1



Agentic
Enterprise

2



Beyond
the Grid

3



Proof
Before Trust

4



Quantum
Utility

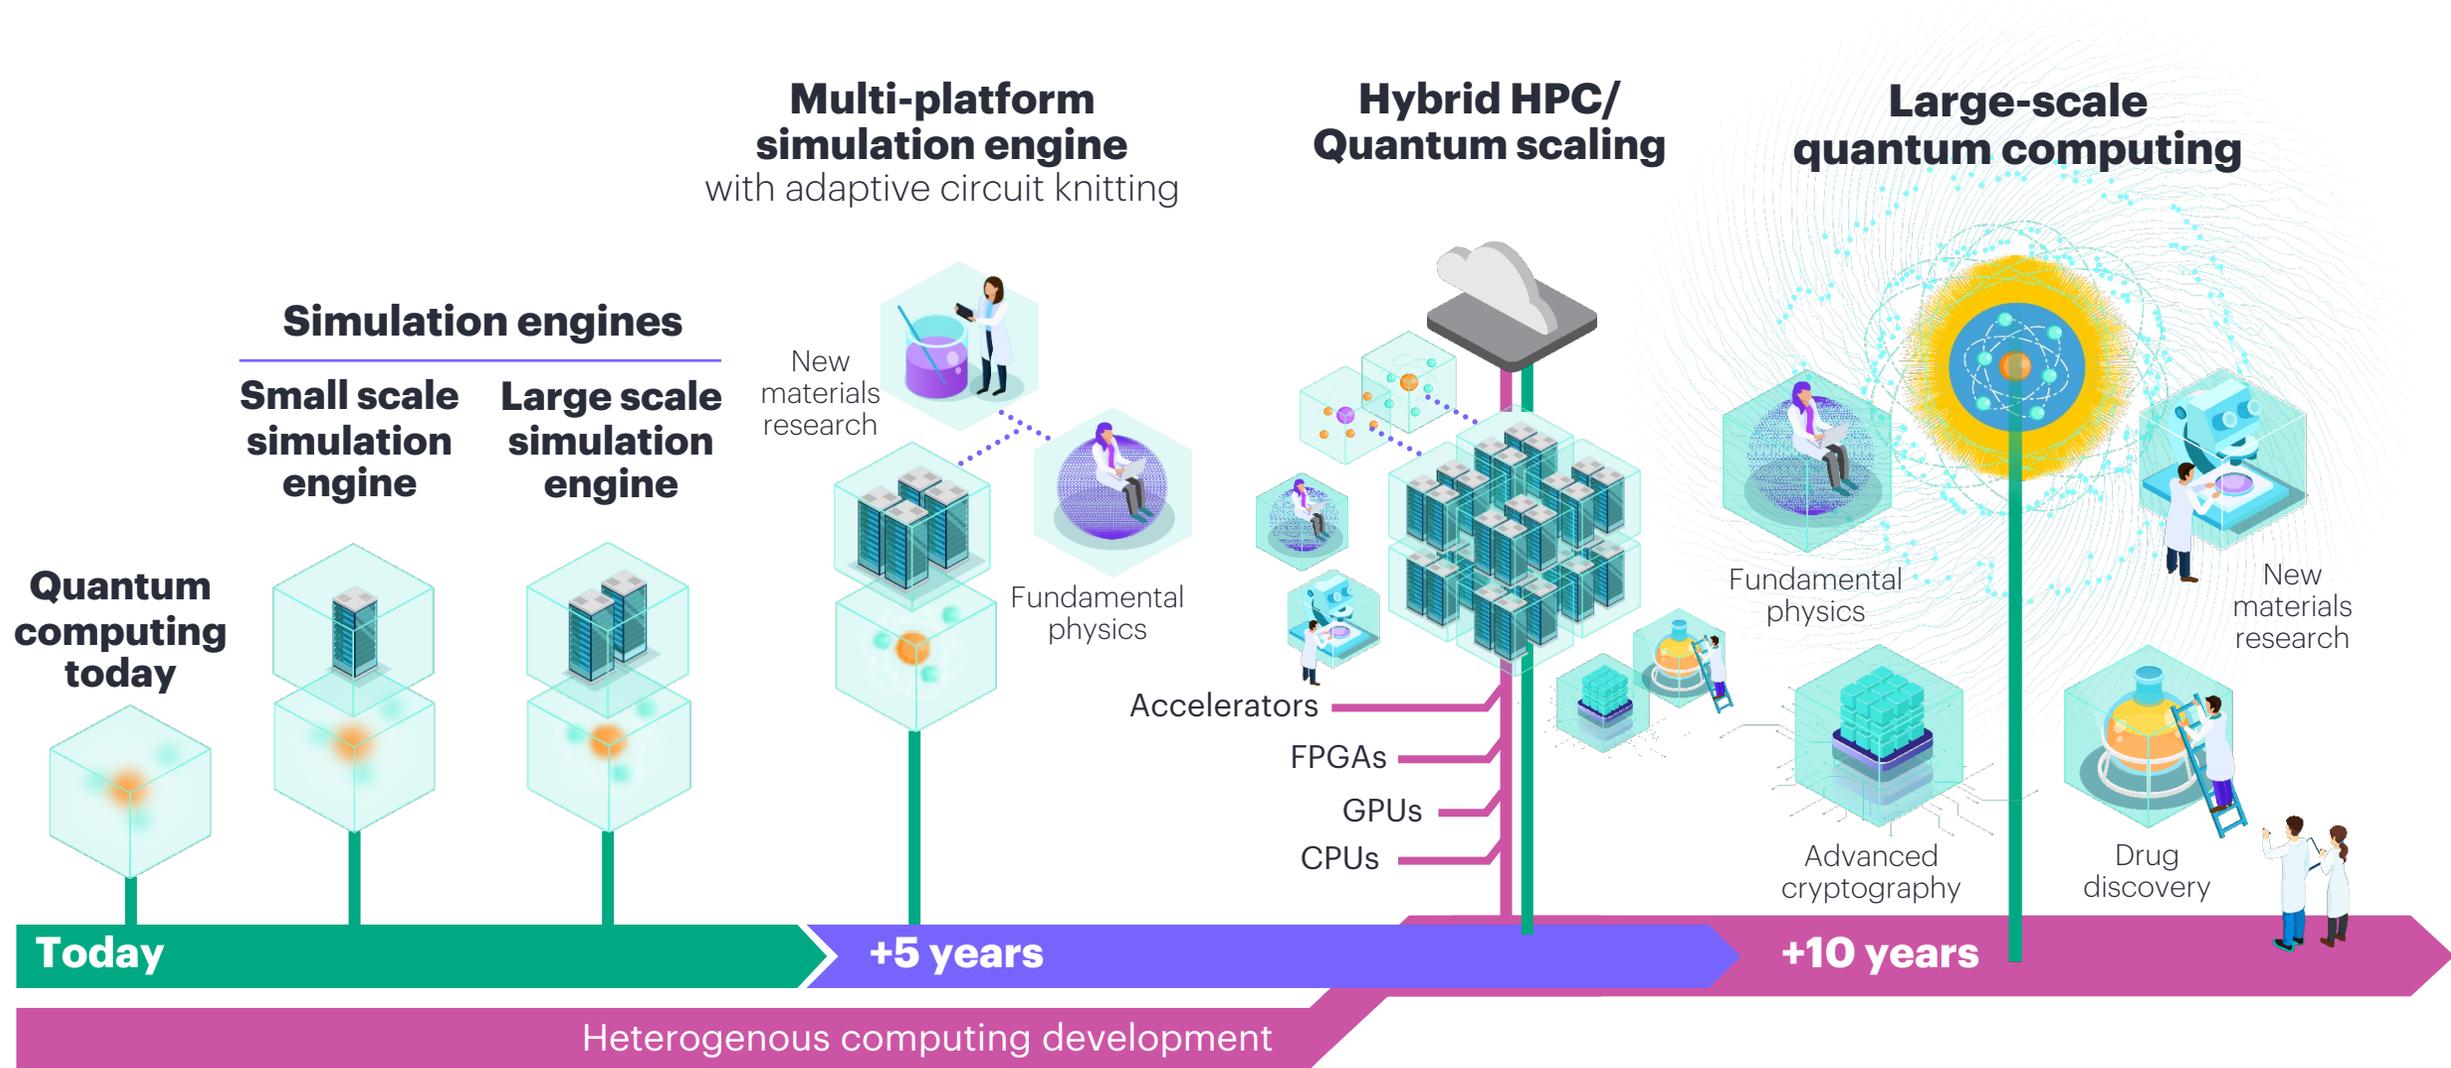
5



AI
After LLMs



The journey to Quantum Computing

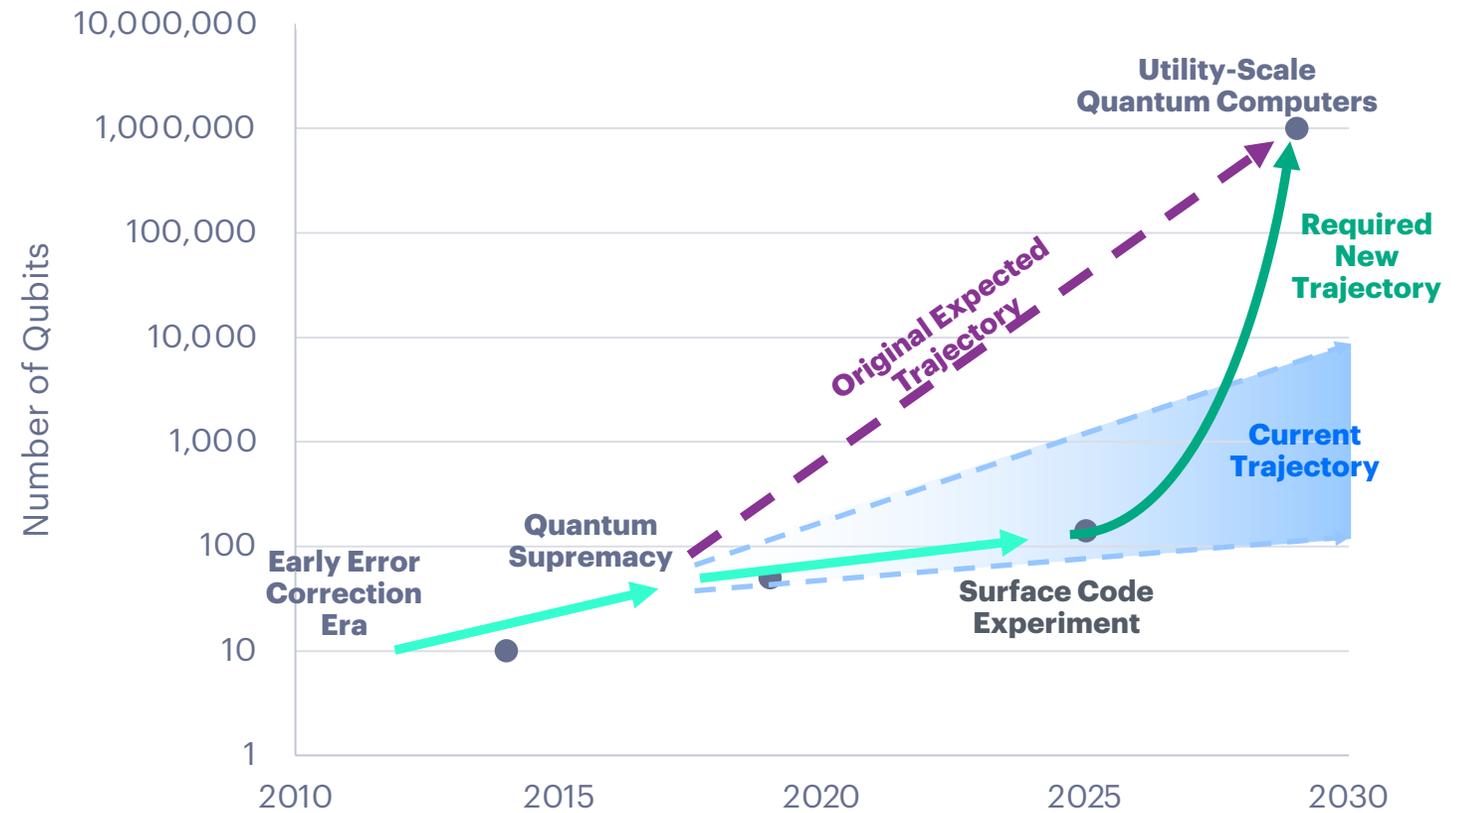


Unlocking the promise of Quantum requires a change in the pace of development

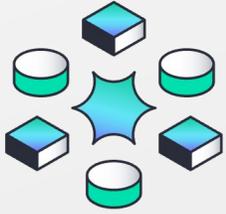
Commentary

- Reaching utility-scale by 2033 requires **increasingly unrealistic scaling trajectories**
- **Many orders of magnitude** separate the qubit counts of today and the counts necessary for tomorrow's utility-scale machines
- **Millions of phys qubits will be required**
Standalone quantum systems will continue struggling
- **Heterogeneous solutions are the answer**
Connect multiple quantum processors with traditional HPC software and hardware

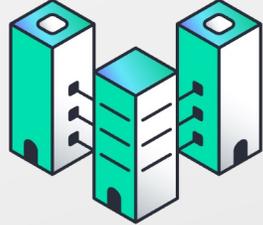
Expectations for Scaling Standalone Quantum Computers



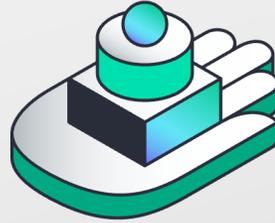
Unlocking Quantum Computing at scale



NETWORKING



HPC/AI



QUANTUM DEV
ENVIRONMENT



HPE Labs

Learn More

Quantum
Scaling Alliance

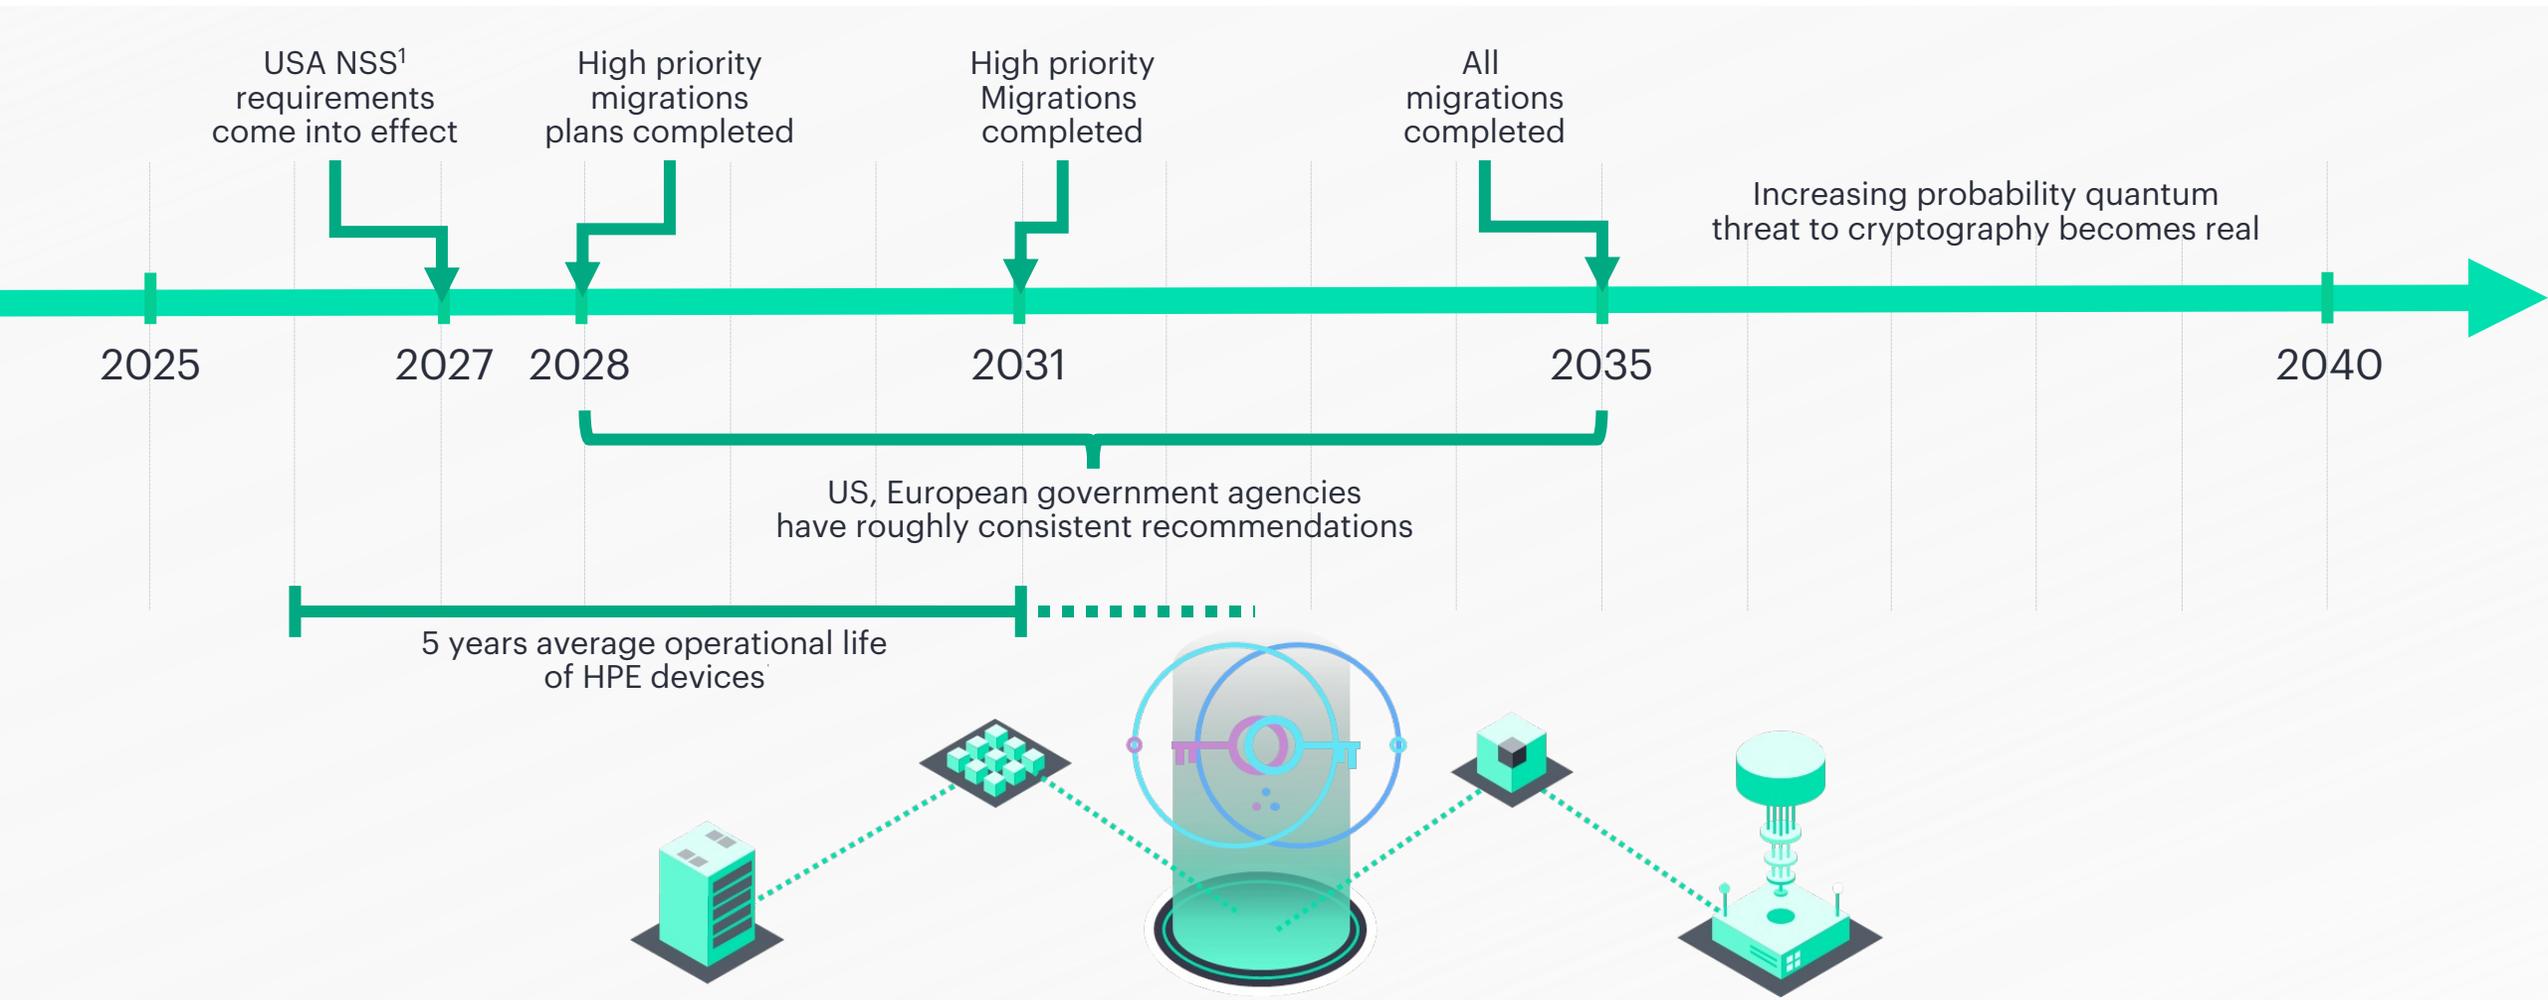


How to Build a Quantum
Supercomputer



Post Quantum Cryptography

Customer PQC migration timeline



The Shape of the World in 2030

1



**Agentic
Enterprise**

2



**Beyond
the Grid**

3



**Proof
Before Trust**

4



**Quantum
Utility**

5



**AI
After LLMs**



Current limitations of LLMs

Practical limitations

- **Limited access to high-quality data**
- **Cost of pre-training, fine-tuning** (RLHF) , and inference
- Limited access to GPU

Computational limitations

- **Quadratic scaling of compute** with number of model parameters in pre-training
- **Combinatorial hardness** and latency in inference (for multimodal data)

Fundamental limitations

- **Trustworthiness**
- Alignment with human/customer values
- **Lack of ability for complex reasoning** (compositional gap persists with scaling)
- Lack of understanding emergence skills

- **Lack of complex reasoning** significantly limits LLM applications to decision making, planning, analytical thinking, scientific reasoning
- **Black-box nature of LLM** limits applications that require guarantees, confidence interval, transparency, security

Why are LLM's limited for complex reasoning problems?

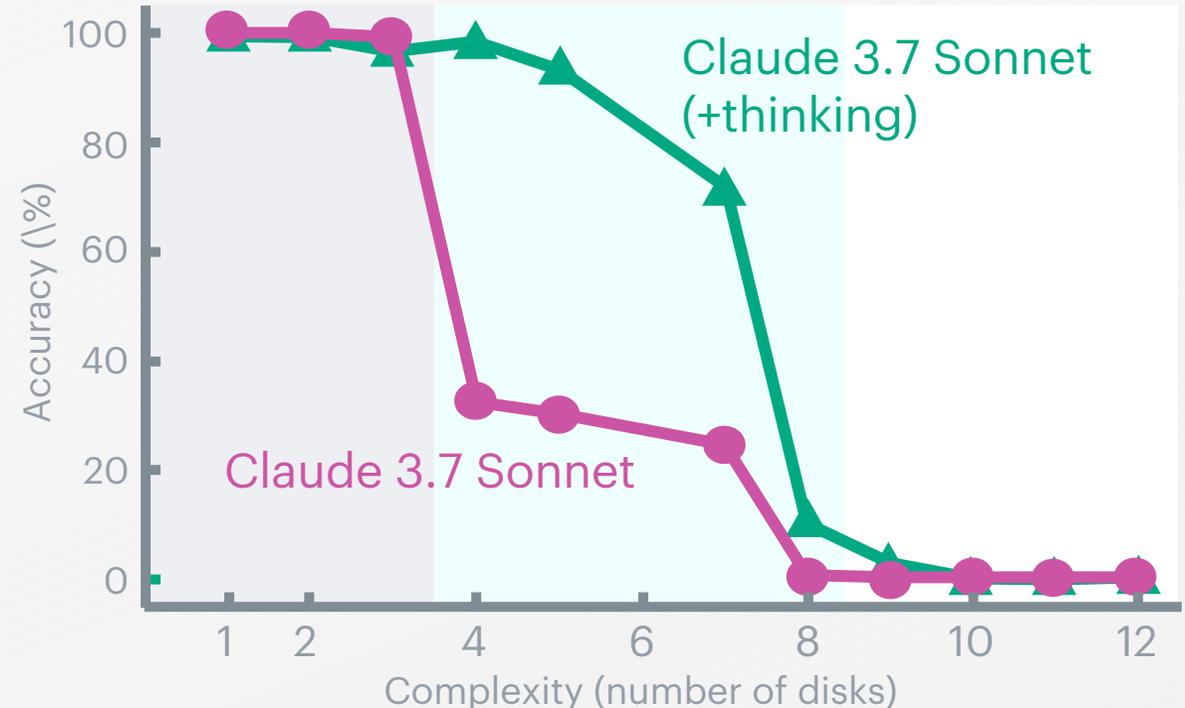
Insights from "The Illusion of Thinking" by Apple (2025)

Three regimes as we increase the complexity of the problem (tower of Hanoi in this case)

For simple problems, LLM (Large Language Model) match or even beats the LRM (Large Reasoning Model based on RL)

In the intermediate, LRM opens a gap in performance and outperforms the LLM

For sufficiently complex tasks both completely fail



Modern LRMs often present an **illusion of reasoning** — fluent explanations without scalable, consistent problem-solving



How are we addressing the limitations of LLMs?

Daniel Kahneman (Nobel laureate in Economic Sciences 2002):

“System 1 operates automatically and quickly, with little or no effort and no sense of voluntary control.”



“System 2 allocates attention to the effortful mental activities that demand it, including **complex computations.**”



How are we addressing the limitations of LLMs?

Augmenting LLMs with system 2 thinking

System 1 =
LLM



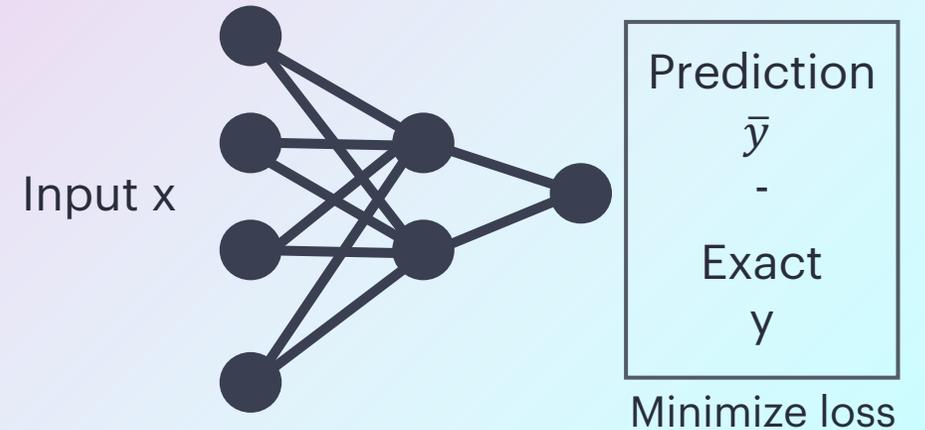
System 2 =
Energy Based Model (EBM)



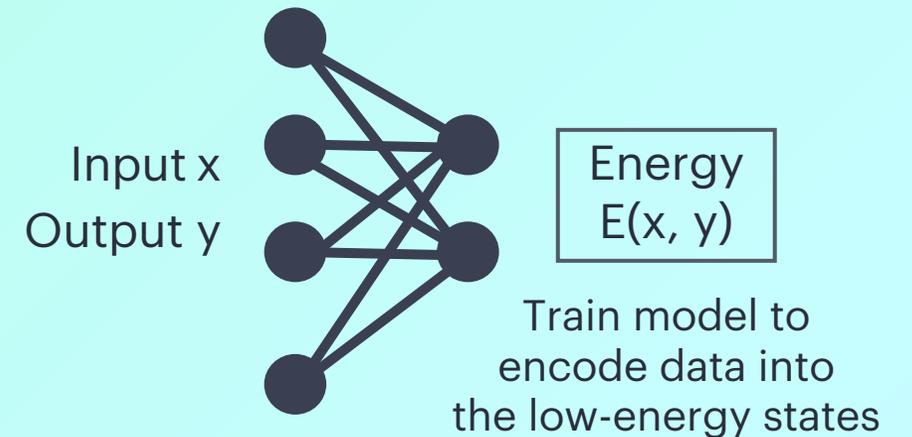
Energy based models Benefits

- **Unified framework** for unsupervised, supervised, or self-supervised learning
- More compact models (high sample quality)
- **Quadratically less compute** than existing LLM as input size increases (linear scaling with tokens)
- **Confidence interval** for generative models (reliability and alignment, similar to Bayesian learning)
- Interpretable/explainable emergence properties
- Directly can be trained or fine-tuned for reasoning and planning tasks
- Applicable to causal learning
- Generally hard to train (focused work at Labs)

Feed-forward model



Energy-based model



Emergent Machine Intelligence

Physics-inspired heterogeneous HPC for scientific computing, complex reasoning, and building world models

Scaling by integrating classical and quantum systems
to harness diverse accelerators that maximize run-time, efficiency, sustainability

Full-stack quantum supercomputer
For efficient quantum simulations and quantum machine learning

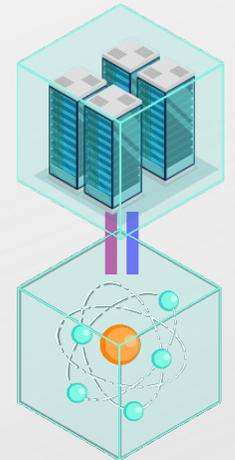
Distributed quantum algorithms and applications
Toward utility scale quantum computing

Hybrid Quantum-Classical Sampling
Solve combinatorial optimization and faster inference

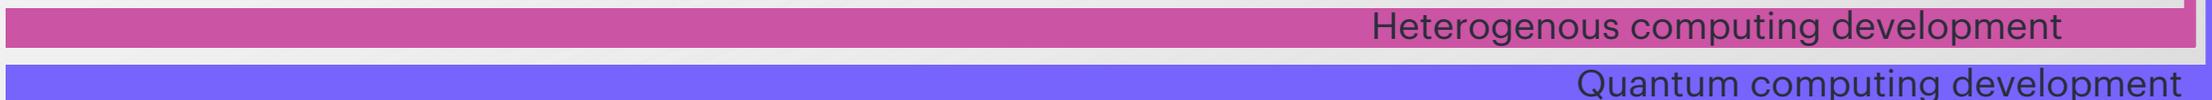
Physics-based generative AI
Augmenting LLMs for complex reasoning



Integration of quantum accelerators



- Accelerators
- FPGAs
- GPUs
- CPUs



Thank You

andrew.wheeler@hpe.com
hpe.com/labs

