BERKELEY LAB

Bringing Science Solutions to the World

~~SCIENTIFIC COMPUTING~~ BEYOND THE EXASCALE ERA

HPC

MODSIM’22 WORKSHOP
SEATTLE, AUGUST 11, 2022

Horst Simon, Berkeley, CA, USA

hdsimon@lbl.gov

Image: OLCF

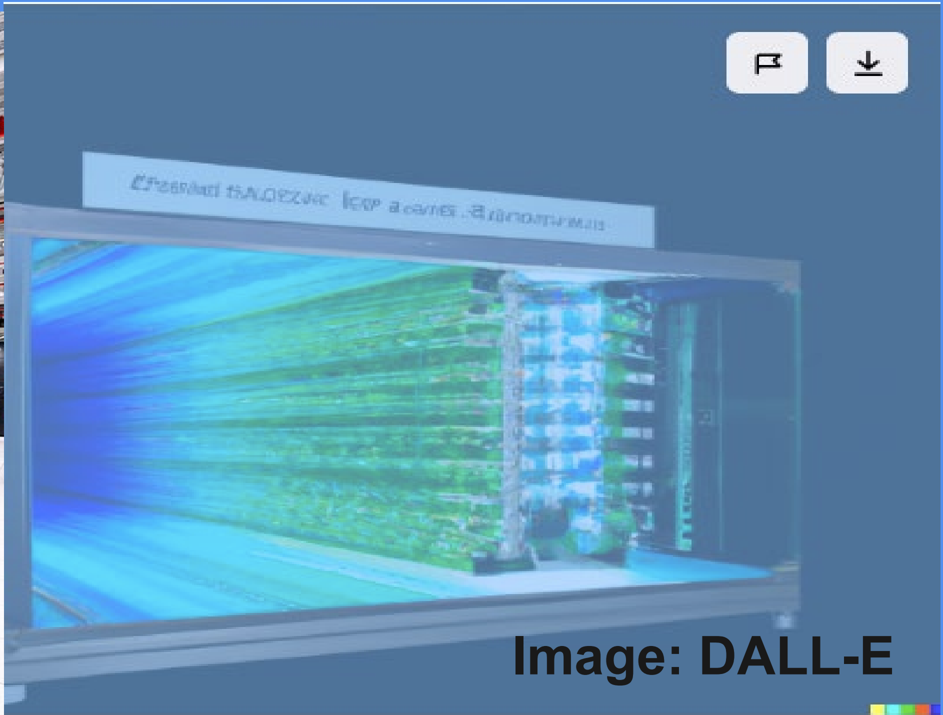
Image: DALL-E

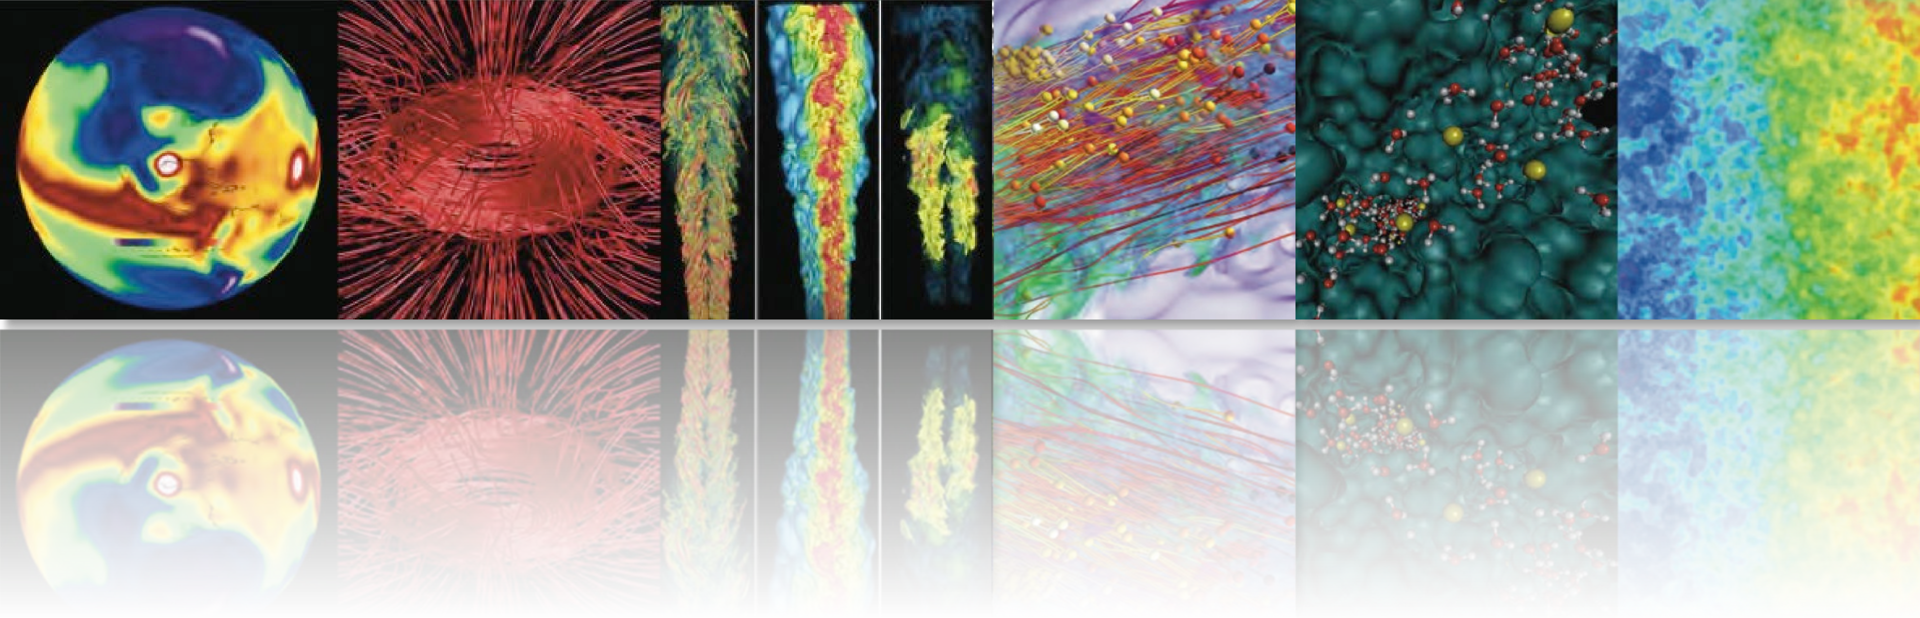

1 Current Market Trends in HPC
(Fragmenting, Domination, Cost)

2 AI/ML for Science

3 Where Will This All Lead To?

BERKELEY LAB

Virtuous Cycle of a General Purpose Technology (semiconductors)

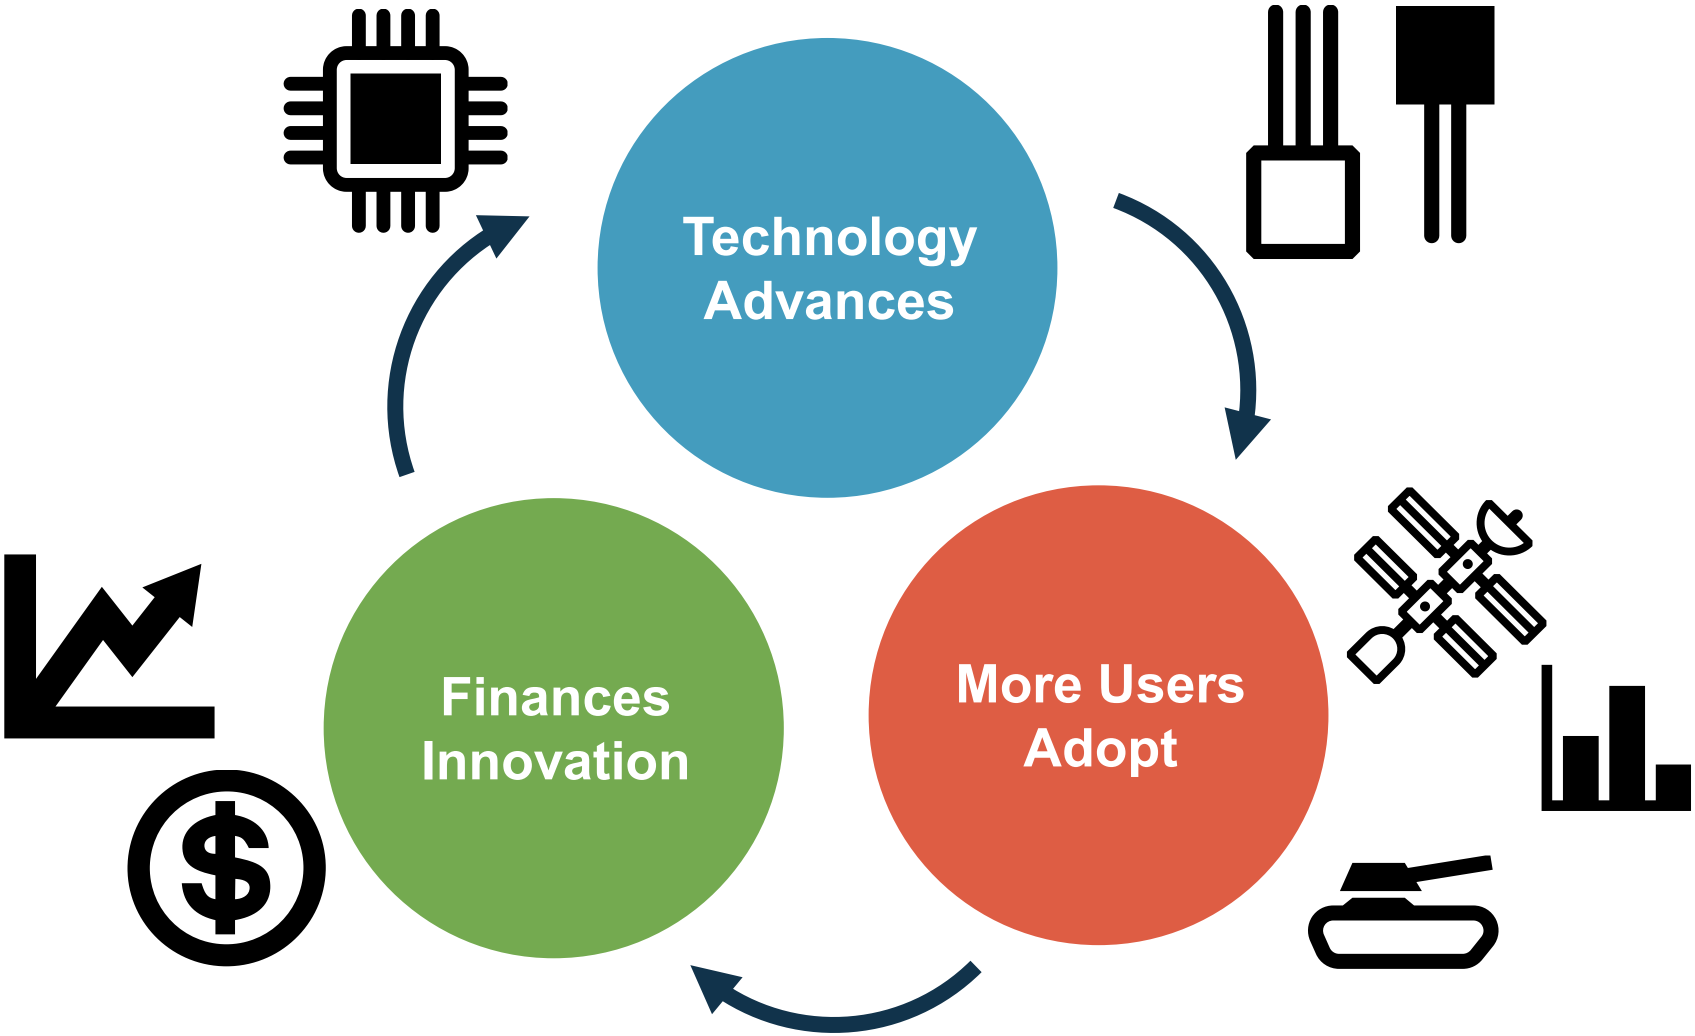

From Neil Thompson, MIT
Lecture at ISC 2019

BERKELEY LAB

Fragmenting of a General Purpose Technology

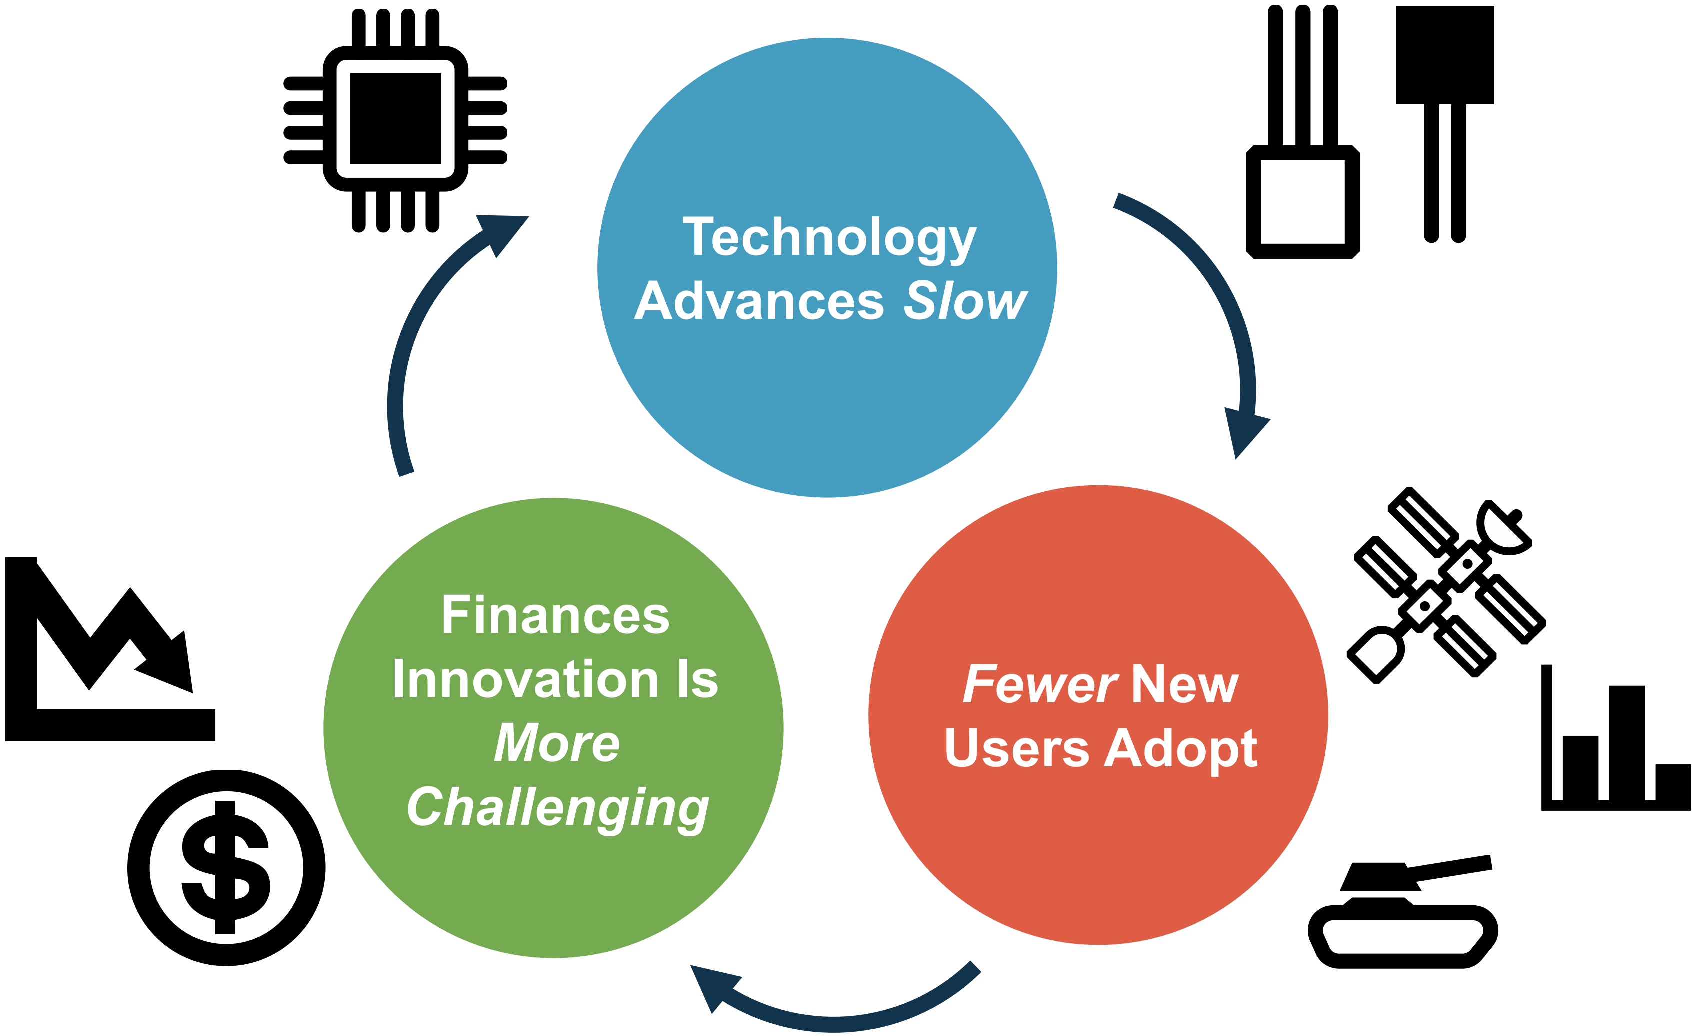

From Neil Thompson, MIT
Lecture at ISC 2019

BERKELEY LAB

Computing is Fragmenting – and so is HPC

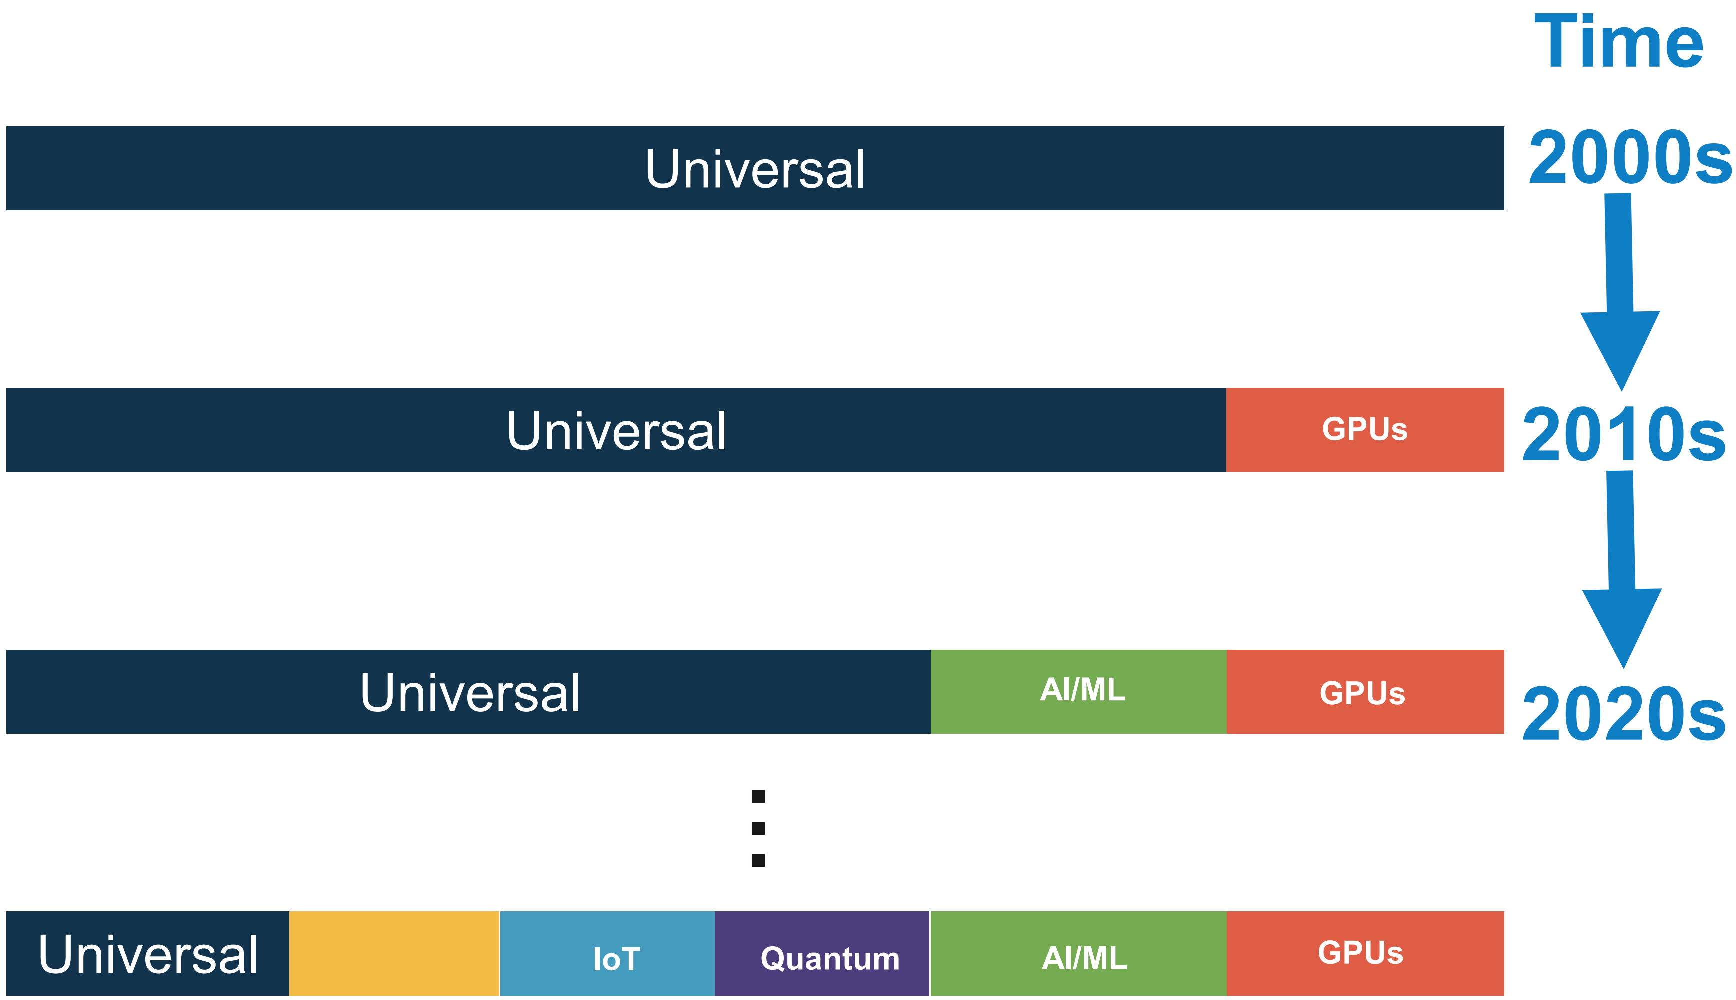

Adapted From Neil Thompson, MIT
Lecture at ISC 2019

BERKELEY LAB

The Cambrian Explosion

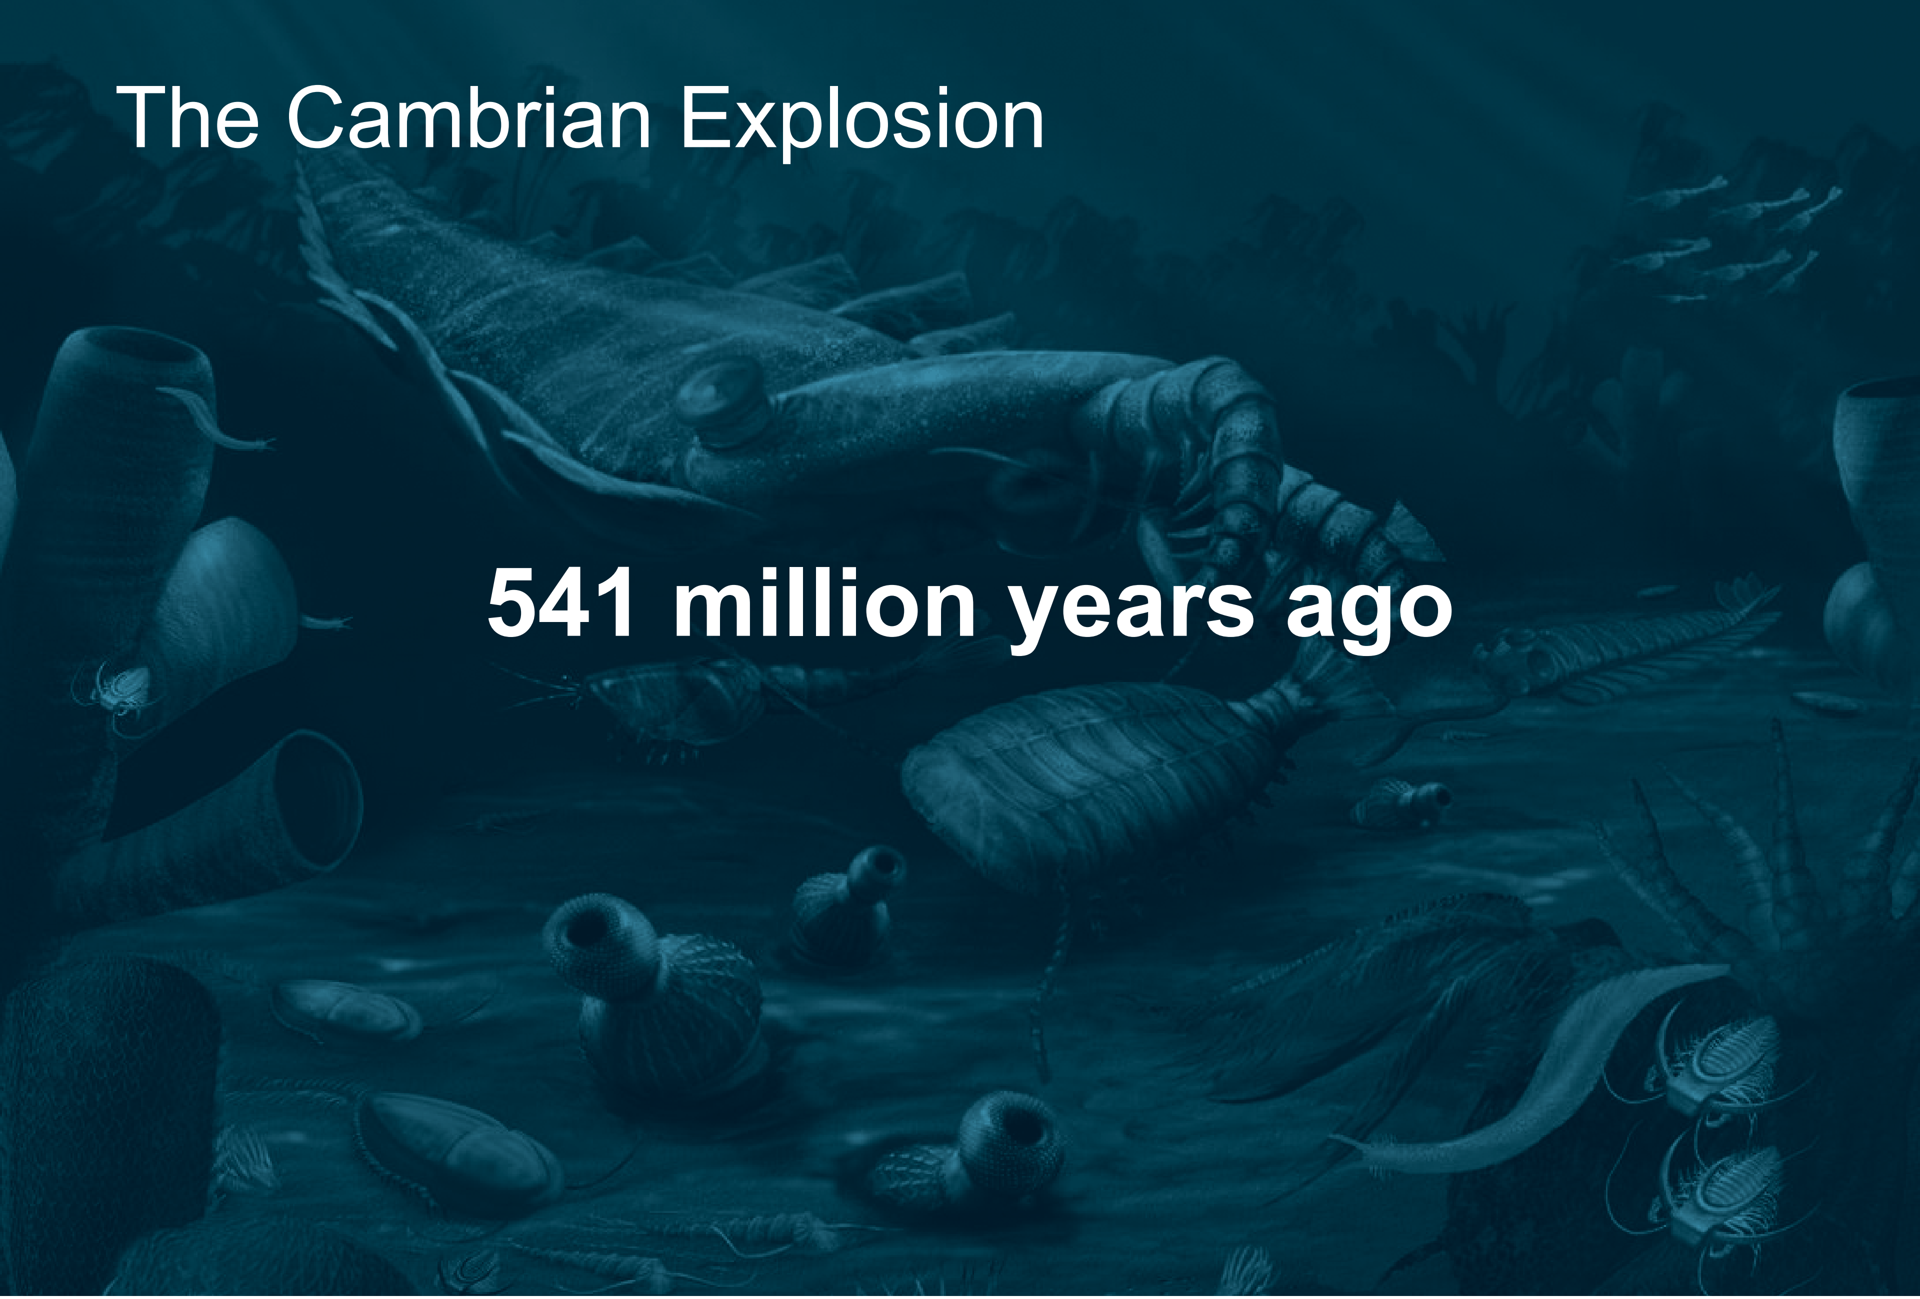

541 million years ago

BERKELEY LAB

The Cambrian Explosion

Most present animal phyla appeared during this period.

The Cambrian Explosion of HPC

Fast Forward to 2022

The Cambrian Explosion of HPC

BERKELEY LAB


Internet of Things
summit
OAK RIDGE
IBM


Internet of Things
summit
OAK RIDGE
IBM
D:Wave
D:Wave

Fragmentation

A wealth of technology choices for HPC

We are living in a golden age of exploration, ….

… but the future is as clear as mud.

BERKELEY LAB

Market Concentration (FAANG+BAT)

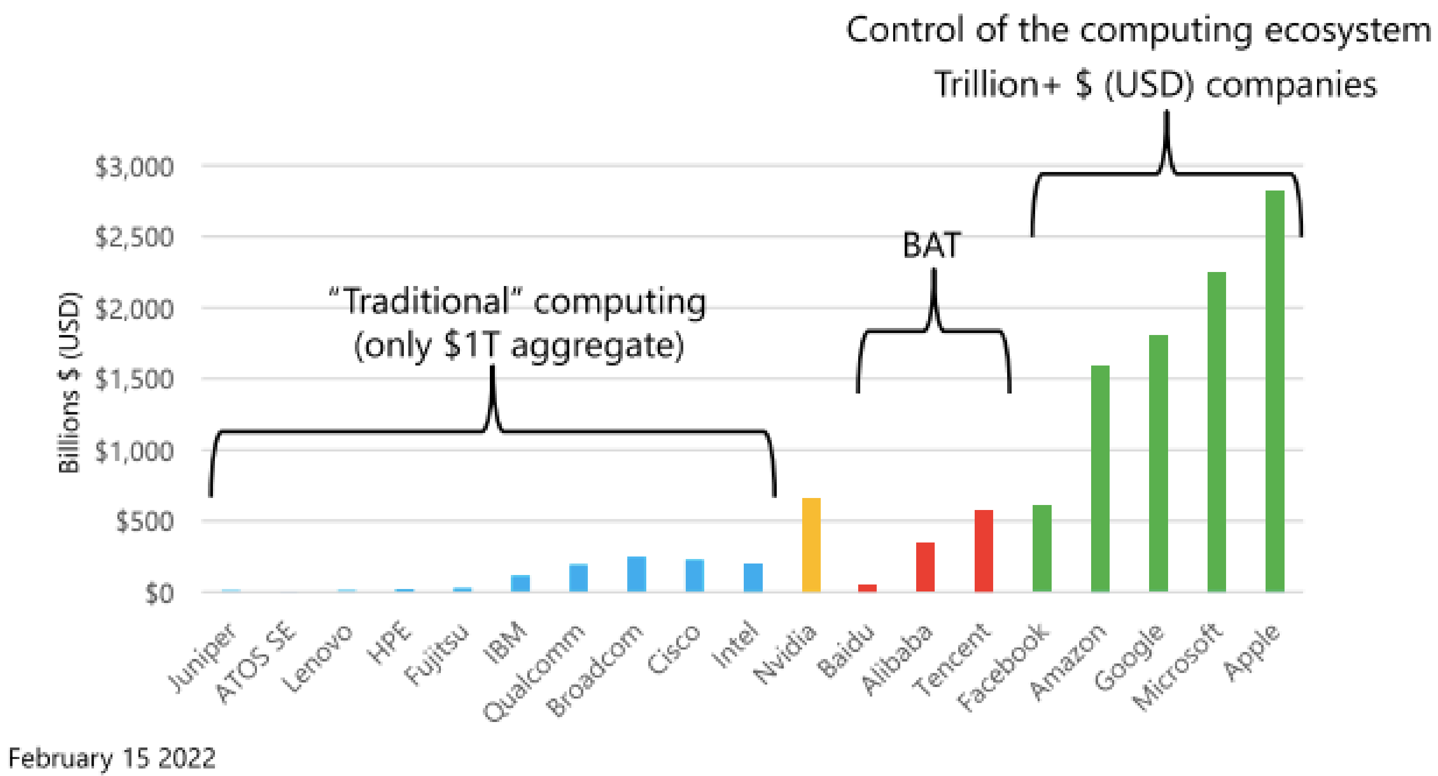

Figure 6: Computing Company Market Capitalization

From Reed, Gannon, and Dongarra https://arxiv.org/abs/2203.02544

Biggest Hyperscale Sites
(http://worldstopdatacenters.com/hyper-scale-data-centers/)

Dwarf HPC sites

Top10 > 400 Ksqft

Biggest Hyperscale Sites

This Top Ten ranks the biggest Hyperscale Sites around the world measured by square footage. These data centers power the need astounding amount of information used for telecommunications, cloud computing, and financial firms. These hyperscale sites are some of the biggest data centers around the world due to the major investments from companies like Google, Facebook, and Microsoft.

#1: Microsoft- Chicago, Illinois

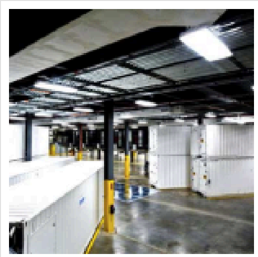

700,000 square feet

Reigning in at number one on our list for biggest data centers is Microsoft. One of the largest data centers ever build so far but also being the most unusual. The two story building uses 40-foot shipping containers to house their web servers and uses the second story for the traditional raised-floor data center space.

Read More . . .

#2: Apple- Maiden, North Carolina

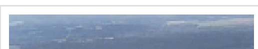

Just barely placing over Google and Microsofts second facility, Apples is number two on our list. The data center itself sits on 183 acres of land that Apple purchased to

BERKELEY LAB

Perspectives on Hyperscalers

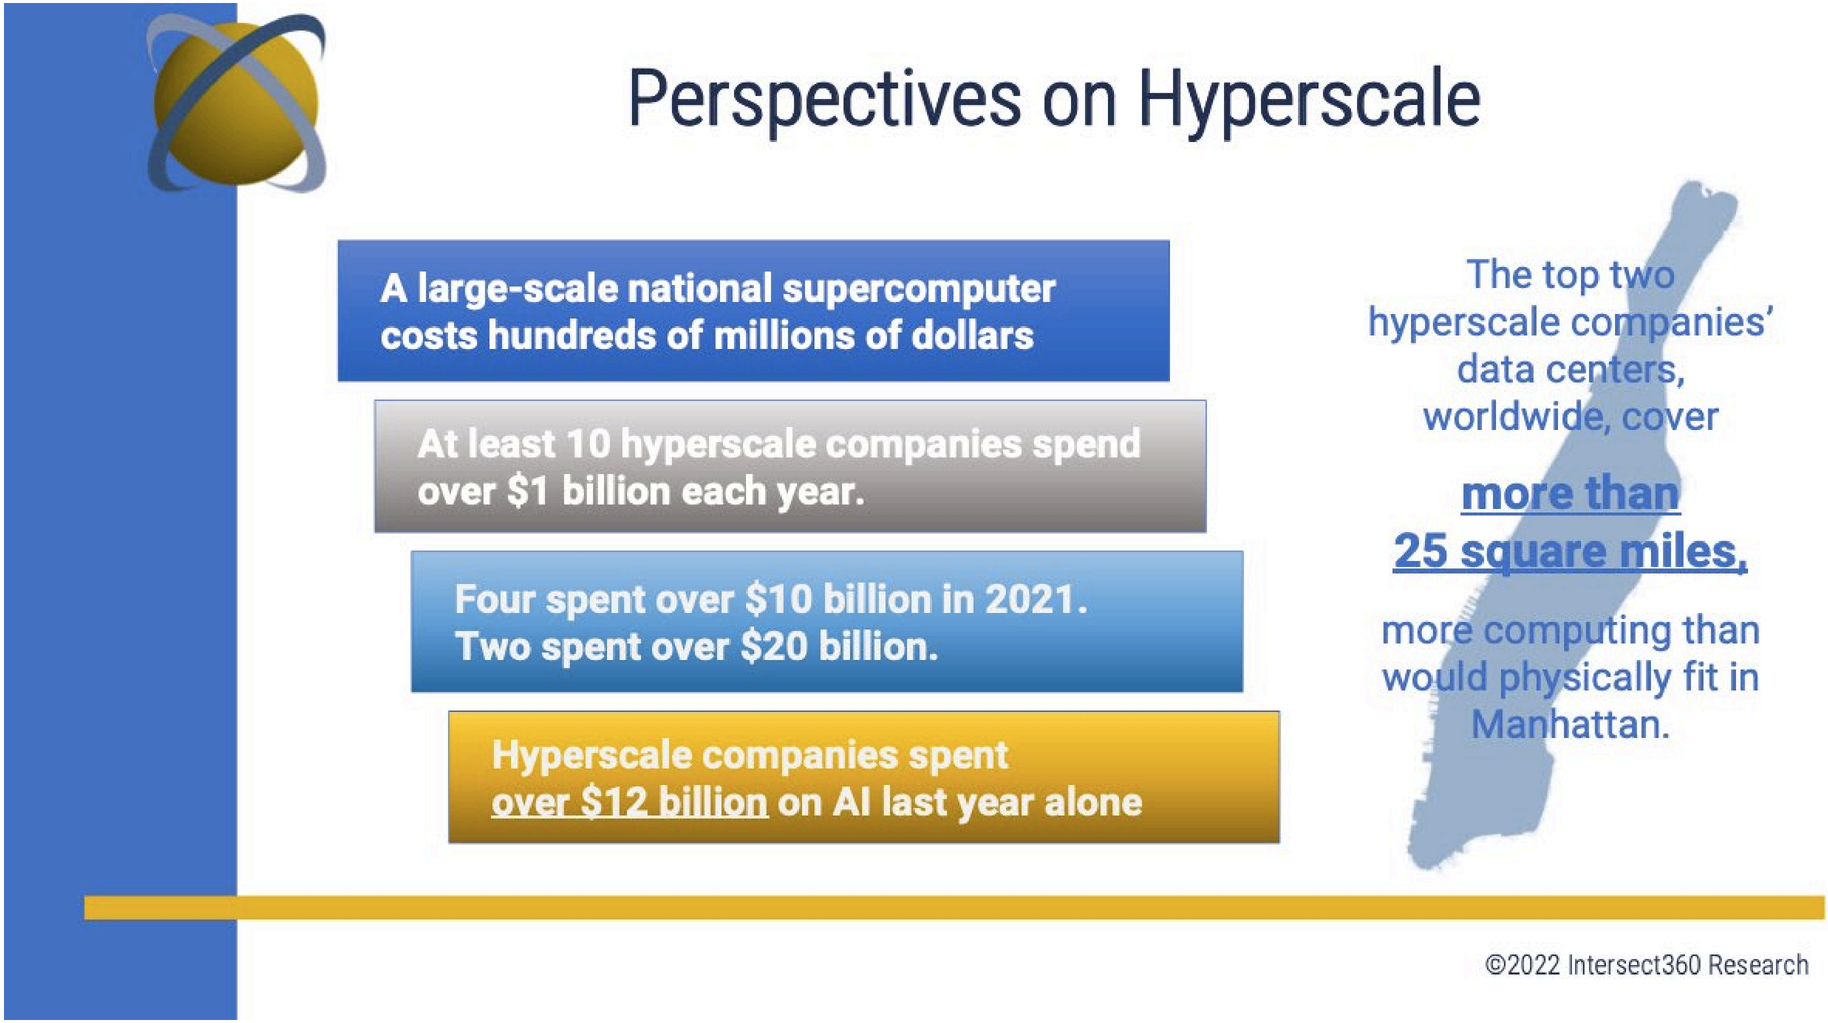

BERKELEY LAB

VENDORS / SYSTEM SHARE

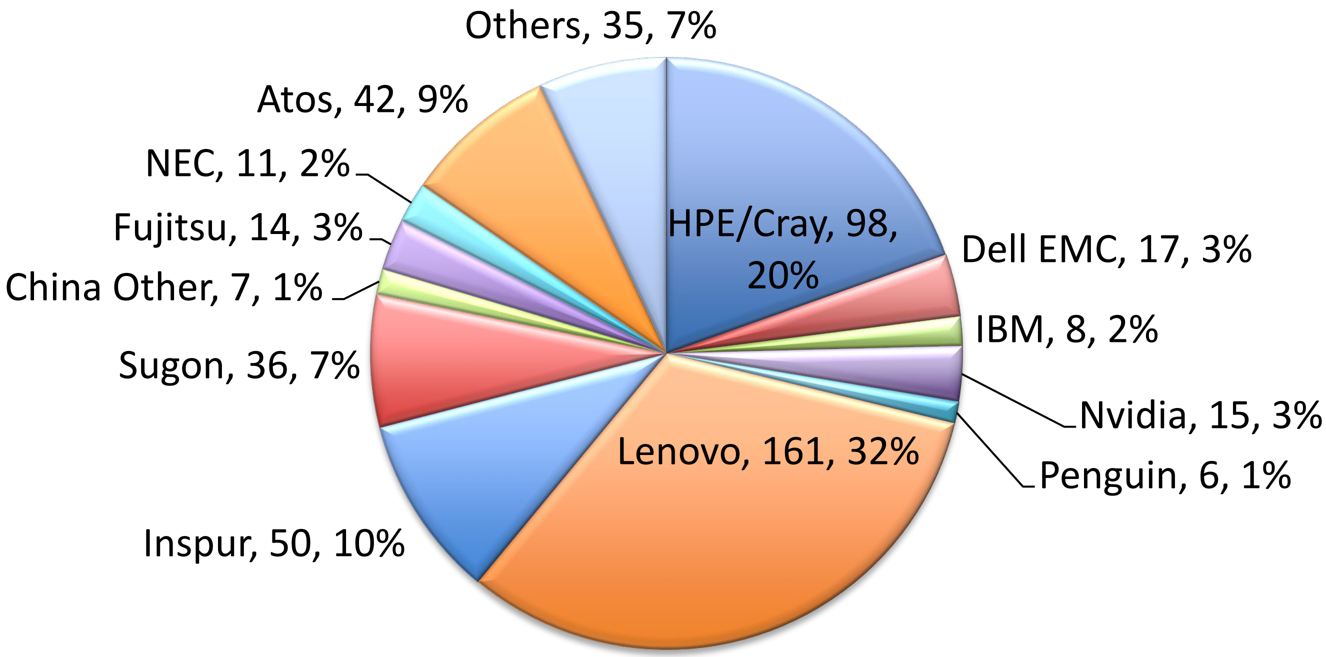

# of systems, % of 500

VENDORS / PERFORMANCE SHARE

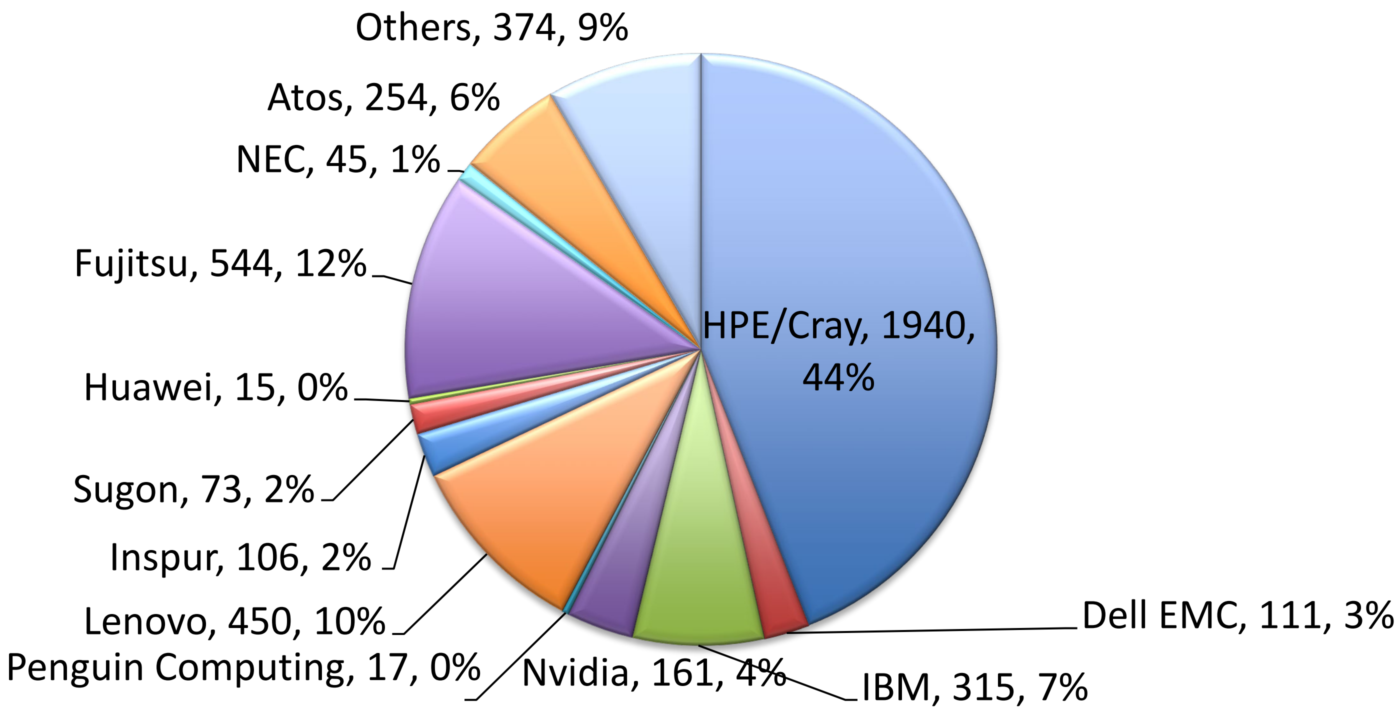

Sum of Pflop/s, % of whole list

Domination by Hyperscalers

Microsoft Datacenter, Quincy

- Financial leverage determines R&D directions, diminishes the influence of HPC on future chip and systems design
- Developing and building their own hardware and software infrastructure at a gigantic scale (TPU, Graviton)
- Determine future directions for Intel, AMD and technology in general
- Loss of HPC systems integrators
- Attract the best talent

Worldwide Race to Exascale (2019)

Projected Pre-Exascale and Exascale Acceptances 2020-2025

Year Accepted	China	EU	Japan	US	Total Installatio	Total Price
2020	1 pre-exascale	1 pre-exascale		1 pre-exascale	3-4	~$750 Million
2021	1 pre-exascale 1 near-exascale	1 pre-exascale	1 (Post K Accepted)	1 pre-exascale	4-5	~$1,900 Million
2022	1 or 2 exascale	1 near-exascale	?	2 exascale	4-5	~$1,700 Million
2023	1 exascale	1 exascale	1 near-exascale ($100 million)	1 or 2 exascale	4	~$1,500 Million
2024	1 exascale	1 exascale	?	2 exascale	4	~$1,400 Million
2025	2 exascale	1 or 2 exascale	1 near-exascale ($100 million)	1 exascale	5-6	~$1,600 Million

Source: Hyperion Research 2019

Power Consumption

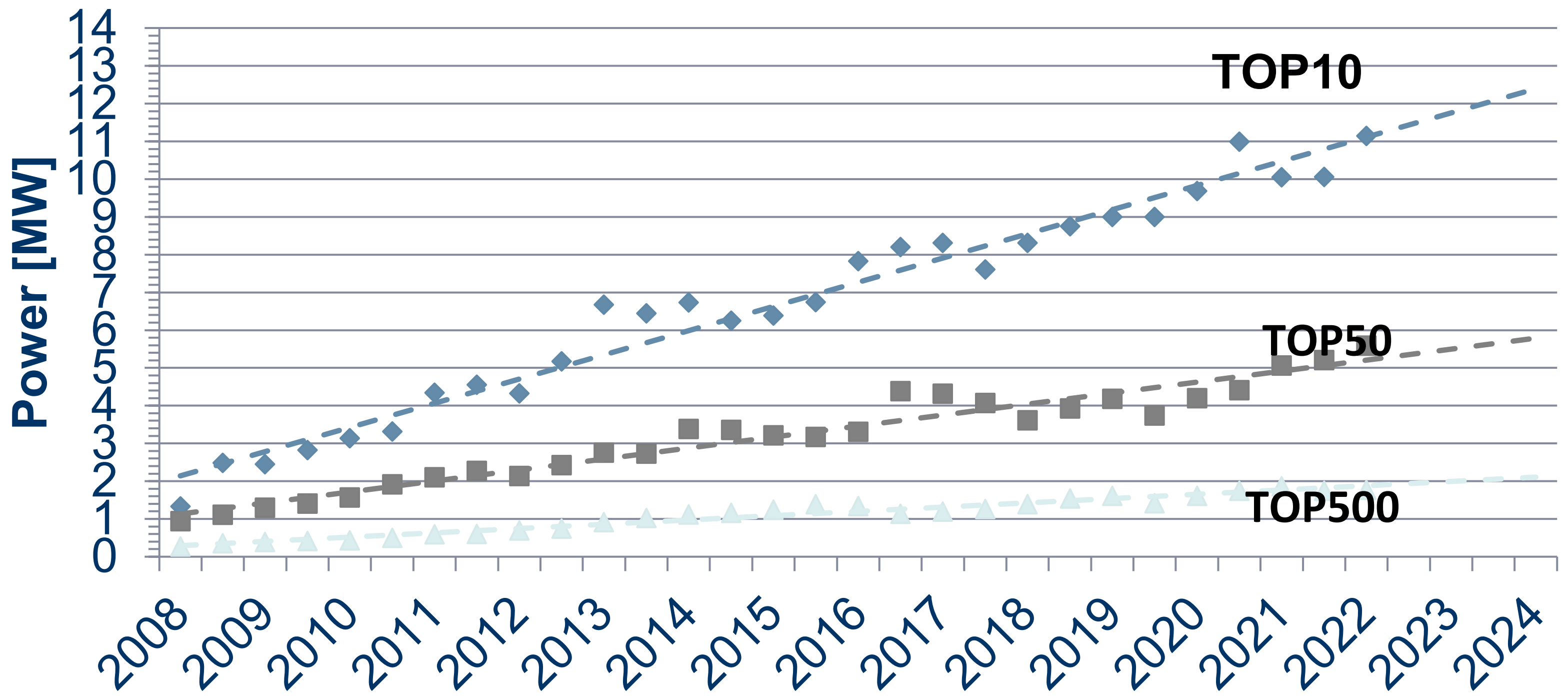

Projected Performance Development (TOP500)

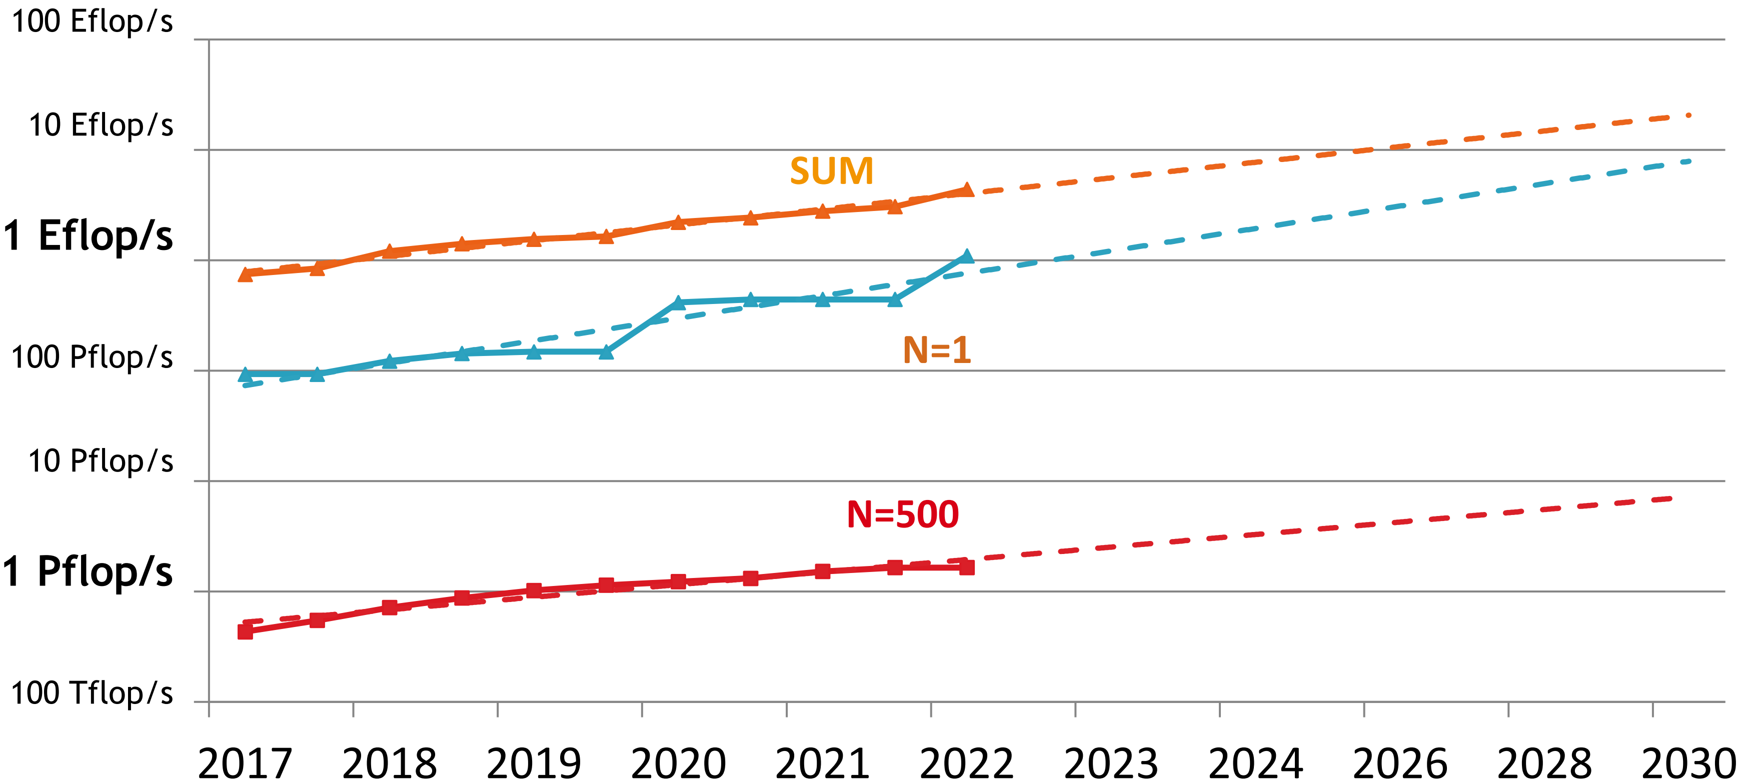

BERKELEY LAB

Costs will Limit Growth of High End HPC Systems

Future TOP10 systems will cost $750M to $1B

Will require of the order of 20MW

I do not expect a 10 Exaflops HPL system before 2030

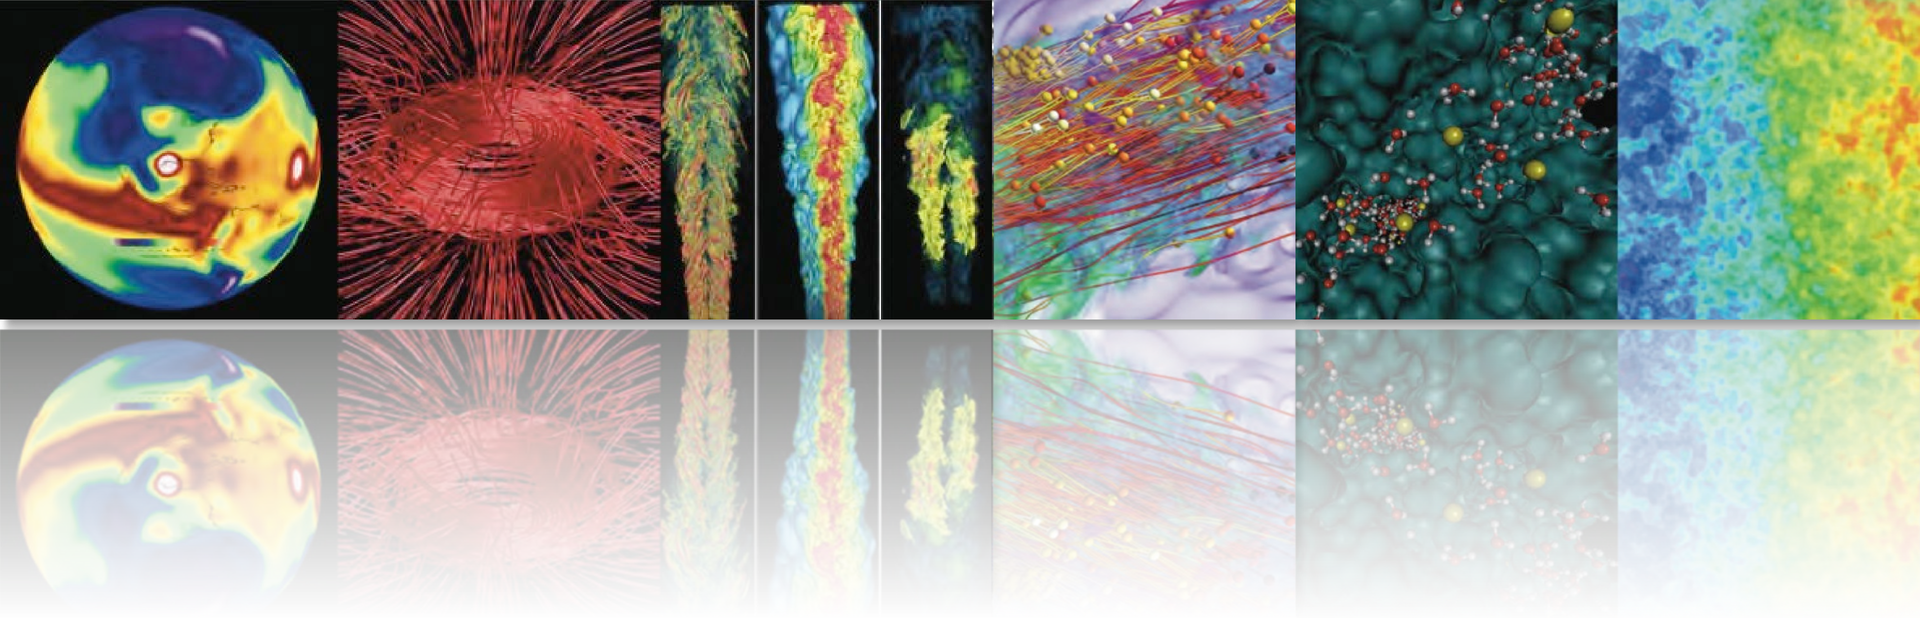

1 Current Market Trends in HPC
(Fragmenting, Domination, Cost)

2 **AI/ML for Science**

3 Where Will This All Lead to?

Computing the Protein Universe

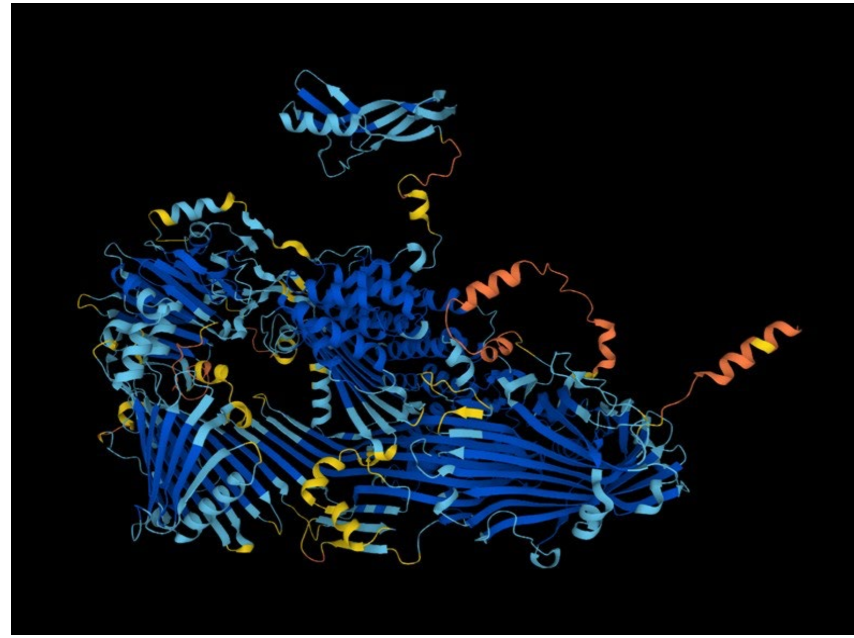

We experience seemingly daily new milestones in machine intelligence:

Alpha Fold DB of Deep Mind (Google) predicts more than 200M protein structures, July 28, 2022.

"AlphaFold is the singular and momentous advance in life science that demonstrates the power of AI."

https://www.nature.com/articles/d41586-022-02083-2

BERKELEY LAB

HPC is essential for progress in large scale AI

300,000x increase from 2011 (AlexNet) to 2018 (AlphaGoZero)

A petaflop/s-day = 10^{15} neural net operations per second for one day, ~= 10^{20} operations

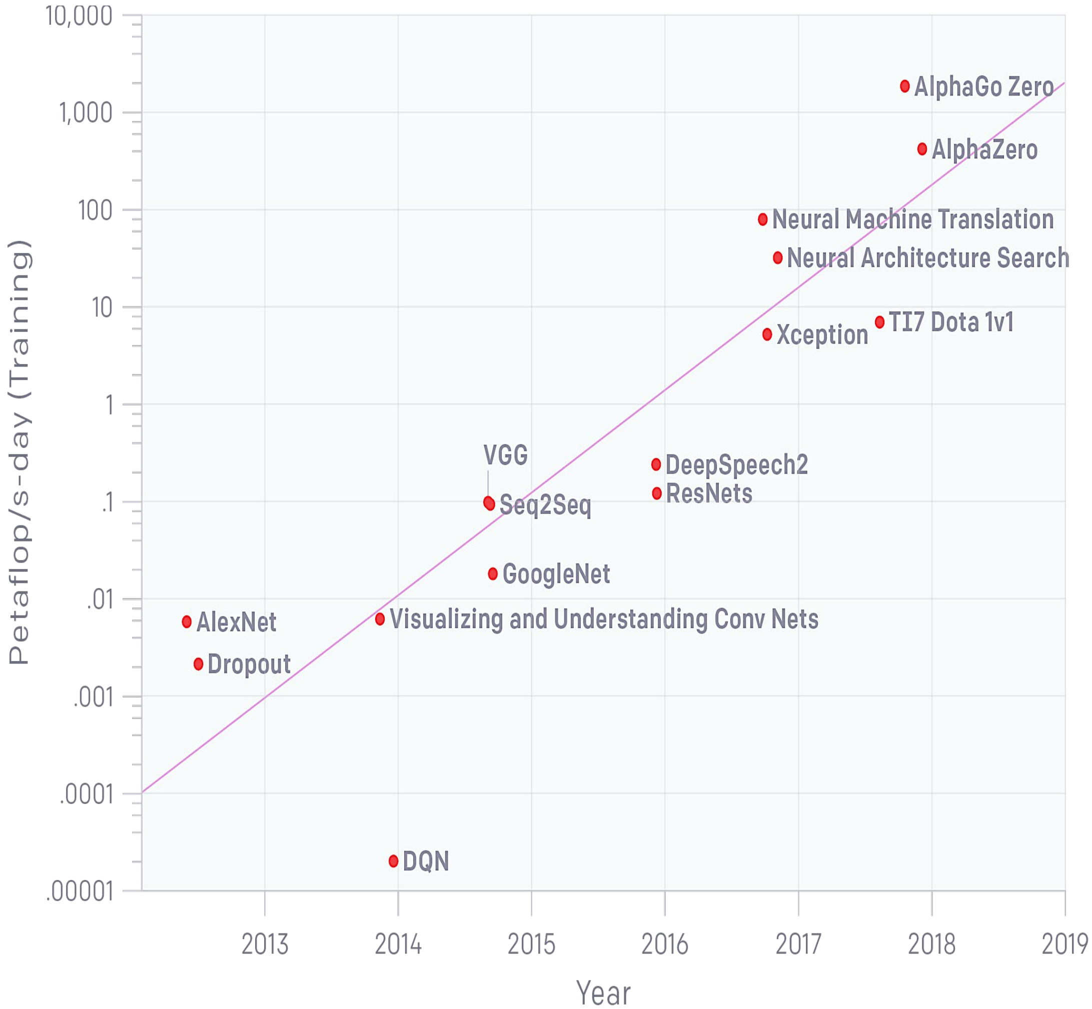

From 2011-2017 the fastest Top500 machine grew < 10x

https://openai.com/blog/ai-and-compute/

BERKELEY LAB

Training Costs of ML Models over Time

(from Sevilla et al. arXIV2202.05924v2)

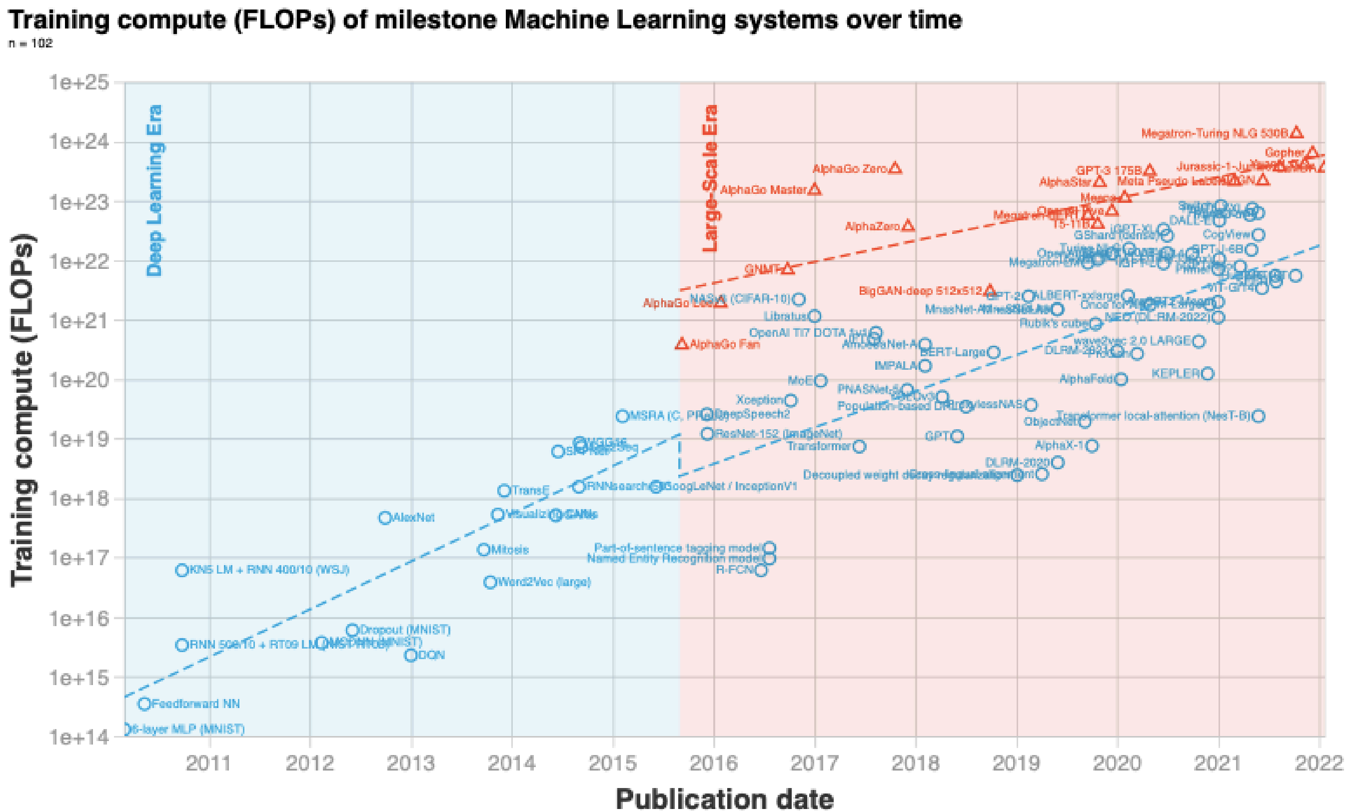

Figure 3: Trends in training compute of n102 milestone ML systems between 2010 and 2022. Notice the emergence of a possible new trend of large-scale models around 2016. The trend in the remaining models stays the same before and after 2016.

These Computing Requirements sparked a high level of Interest in Architectures and Systems for ML

- Explosion of VC interest in chip start-ups
- More than 45 start-ups are working on chips that can power tasks like speech and self-driving cars
- US VC are investing >$4B in start-ups
- SambaNova ($1.2B), Groq ($400M), Cerebras ($250M) are large players in this field


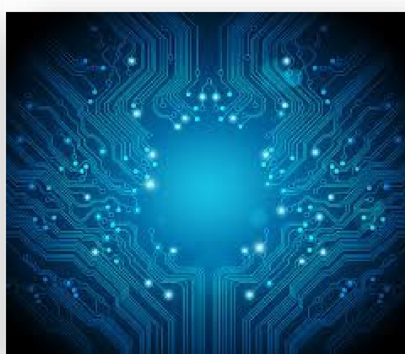
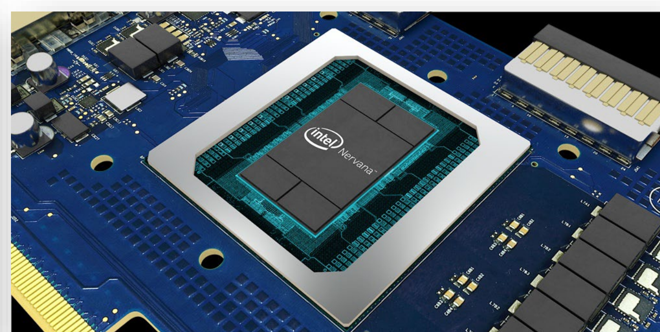

Proliferation of New Companies and Architectures

AI Chip Landscape

More on https://basicmi.github.io/AI-Chip/

Tech Giants/Systems

- Google
- Microsoft
- aws
- IBM
- facebook
- HUAWEI
- Baidu 百度
- Alibaba Group 阿里巴巴集团
- FUJITSU
- NOKIA
- TOSHIBA
- Hewlett Packard Enterprise
- DELL

IC Vender/Fabless

- intel
- SAMSUNG
- NVIDIA
- QUALCOMM
- AMD
- NXP
- ST
- XILINX
- MEDIATEK
- BROADCOM
- MARVELL
- Rockchip 瑞芯微电子

IP/Design Service

- arm
- synopsys
- Imagination
- cadence
- CEVA
- VeriSilicon
- SiFive
- ARTERIS IP
- alchip
- GUC
- FARADAY
- eSilicon

Startup in China

- Cambricon 寒武纪科技
- 地平线 Horizon Robotics
- BITMAIN
- intellifusion 云天励飞
- ChipIntelli
- Think Force
- Canaan
- 云知声 Unisound
- Rokid
- AISPEECH 思必驰 专注人性化的智能语音
- 肇观电子 NextVPU
- Enflame
- 亿智科技
- 清微智能 TSING MICRO

Startup Worldwide

- cerebras
- WAVE COMPUTING
- Graphcore
- habana
- thinci
- SambaNova SYSTEMS
- KALRAY
- LIGHTELLIGENCE
- HAILO Empowering Intelligence
- Esperanto TECHNOLOGIES
- Tenstorrent
- MYTHIC
- Preferred Networks
- brainchip
- PEZY Computing
- GREENWAVES TECHNOLOGIES
- AIMOTIVE
- KONIKU
- Tachyum
- flexlogix
- SYNTIANT
- gyrfalcon technology
- NOVUMIND

扫码访问AI芯片文章

Compiler

- GLOW
- tvm
- NVIDIA TensorRT
- ONNC
- nGraph Compiler stack (Beta)
- plaidML

Benchmarks

- MLPerf
- AI - Benchmark
- AI Matrix.
- AIIA

ALCF AI Testbeds

https://www.alcf.anl.gov/alcf-ai-testbed

Cerebras (CS-2)

SambaNova

Graphcore

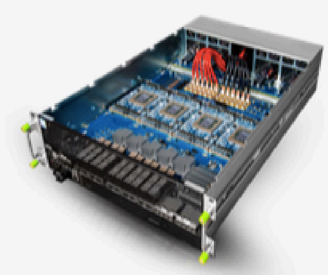

Habana

Groq

- Infrastructure of next-generation machines with hardware accelerators customized for artificial intelligence (AI) applications.
- Provide a platform to evaluate usability and performance of machine learning based HPC applications running on these accelerators.
- The goal is to better understand how to integrate AI accelerators with ALCF's existing and upcoming supercomputers to accelerate science insights

From Venkat Vishwanath, ANL, Venkat@anl.gov

Machine Learning

https://xkcd.com/1838

BERKELEY LAB

Scientific Machine Learning – DOE ASCR BRN Workshop 2019

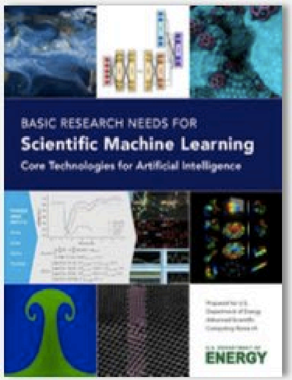

Scientific Machine Learning

What are the opportunities and challenges of machine learning in complex applications across science, engineering, and medicine?

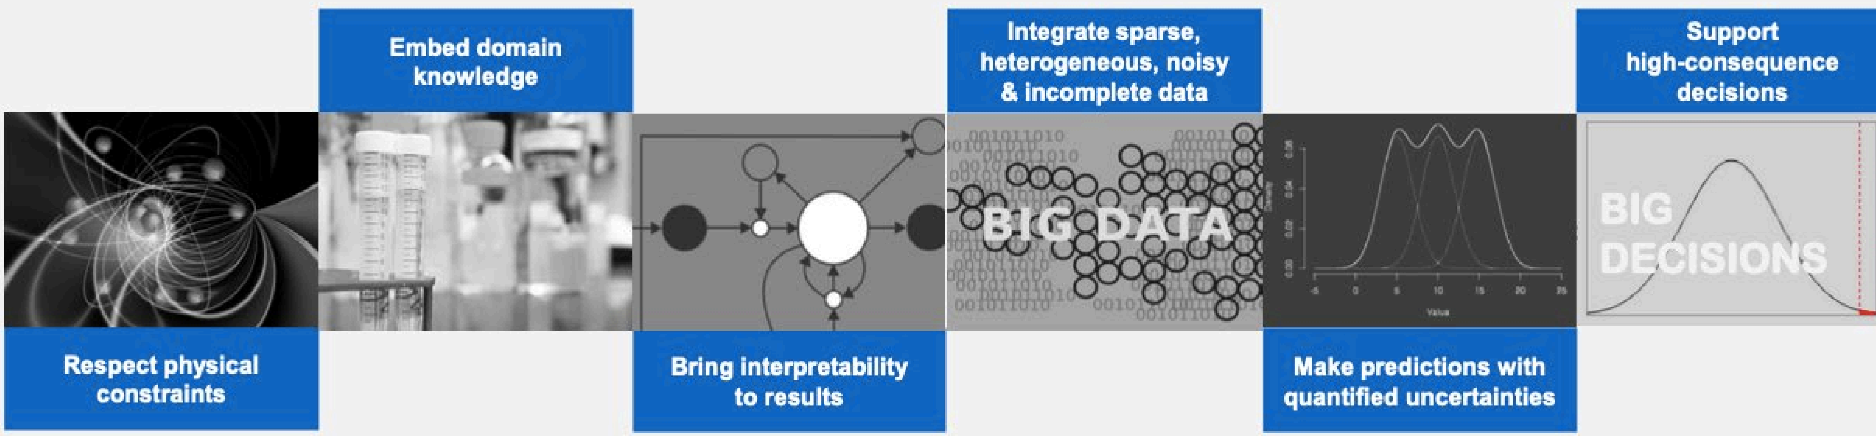

From Karen Wilcox, UT Austin

AI for Science – What's Next After Exascale

- Over 1,300 scientists participated in four town halls during the summer/fall of 2019
- Research Opportunities in AI
  - Biology, Chemistry, Materials,
  - Climate, Physics, Energy, Cosmology
  - Mathematics and Foundations
  - Data Life Cycle
  - Software Infrastructure
  - Hardware for AI
  - Integration with Scientific Facilities
- Modeled after the Exascale Series in 2007
- ASCAC subcommittee Report Sept 2020

https://www.anl.gov/ai-for-science-report

From Rick Stevens, AI4SS workshop

BERKELEY LAB

AI for Science Example Foundation Building Blocks

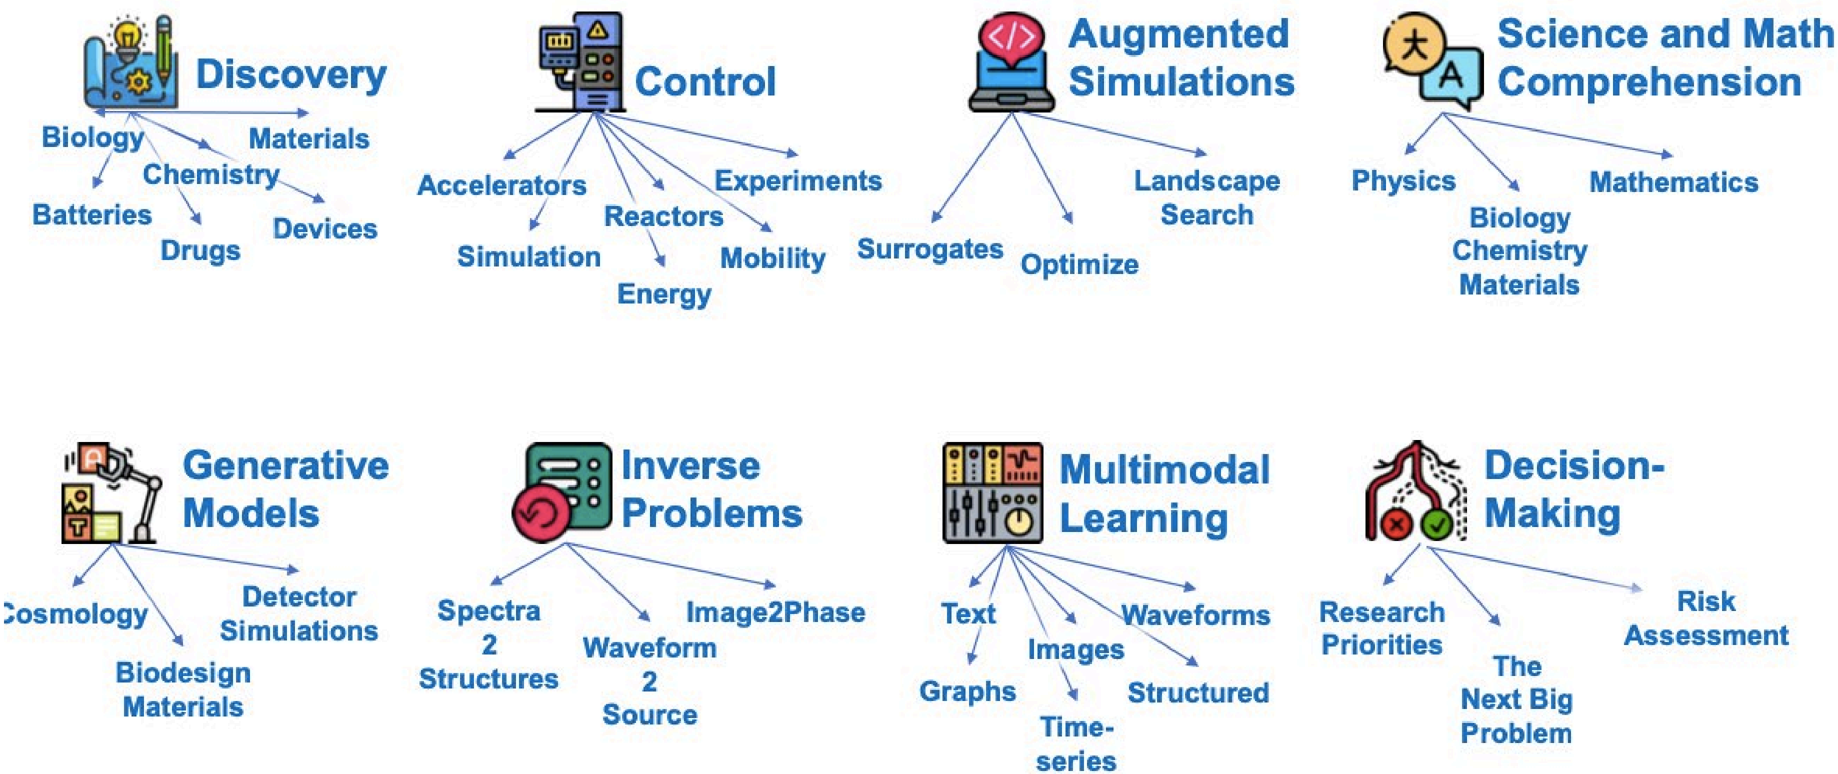

From Rick Stevens, AI4SS workshop

AI for Science "SuperGeneral" Models

AI Enabled
Design Workflows
(what to make)

→ ...materials, polymers, organisms...

AI Enabled
Experimental Workflows
(how to make it)

→ ...self-driving labs, synthesis search...

AI Enabled
Scientific Comprehension
(what it means)

→
- data Sets
- literature
- science "news"
- strategy

→ *Cleaned*
Updated
Annotated
Aggregated
Interpreted
→ Insight?

From Rick Stevens, AI4SS workshop

Summer 2022 Workshop Topics

June, July, August

Argonne, Berkeley, Oak Ridge, Livermore, Los Alamos, Sandia

AI for Advanced Properties Inference and Inverse Design	AI and Robotics for Autonomous Discovery	AI Based Surrogates for HPC
AI for Programming and Ssoftware Engineering	AI for Prediction and Control of Complex Engineered Systems	Foundation AI for Scientific Knowledge Discovery, Integration and Synthesis

From Rick Stevens, AI4SS workshop

Summary on AI for Science

Several workshop reports list research directions for AI/ML and how it can accelerate modeling and simulation:

- surrogate models
- digital twins
- physically constrained models

National Labs bring together team science, large scale collaborations, and exascale computing experience

- expect AI4SS to develop “grand challenge” type problems that will set the agenda for the next five years and that can be approached by large scale models

BERKELEY LAB

SciML Capabilities

Machine Learning for Advanced Scientific Computing Research

Data-intensive scientific inference & data analysis	ML methods for multimodal data in situ data analysis with ML ML to optimally guide data acquisition ⋮
ML-enhanced modeling & sim ML-hybrid algorithms and models for better scientific computing tools	ML-enabled adaptive algorithms ML parameter tuning ML-based multiscale surrogate models ⋮
Intelligent automation & decision support automated decision support, adaptivity, resilience, control	exploration of decision space with ML ML-based resource mgt & control optimal decisions for complex systems ⋮

SciML Foundations

Machine Learning for Advanced Scientific Computing Research

Domain-aware leveraging & respecting scientific domain knowledge	physical principles & symmetries physics-informed priors structure-exploiting models ⋮
Interpretable explainable & understandable results	model selection exploiting structure in high-dim data uncertainty quantification + ML ⋮
Robust stable, well-posed & reliable formulations	probabilistic modeling in ML quantifying well-posedness reliable hyperparameter estimation ⋮

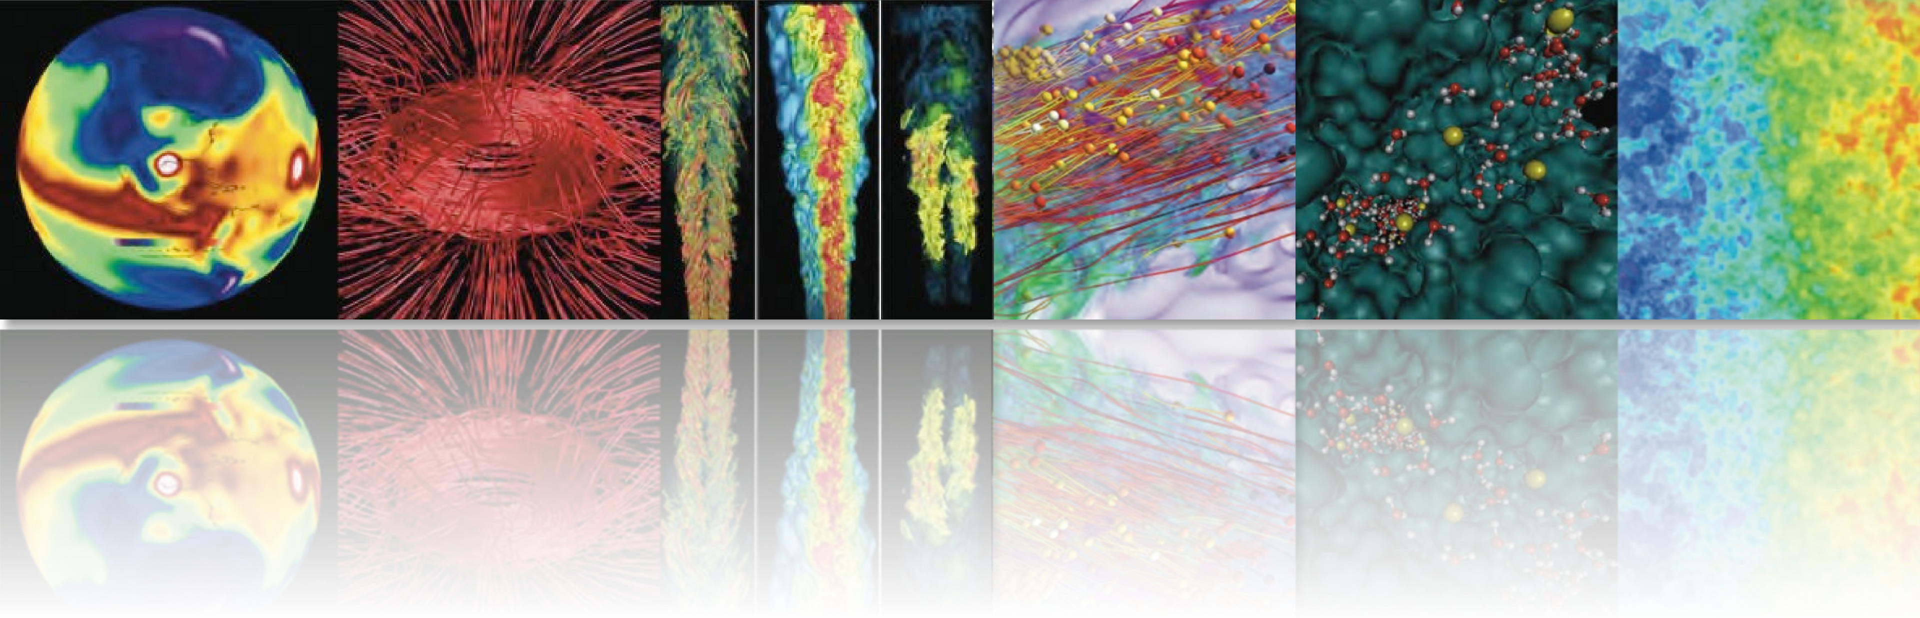

1 Current Market Trends in HPC
(Fragmenting, Domination, Cost)

2 AI/ML for Science

3 Where Will This All Lead to?

BERKELEY LAB

Topics that I did not discuss … (but have opinions about)

Quantum computing

End of Moore's Law

Semiconductors and supply chain

The growth of experimental data and IoT

BERKELEY LAB

A changing computing landscape challenges us to think differently about the HPC computing environment

Growth of experimental and observational data and the need for interactive feedback through real-time data analysis and simulation and modeling

DESI

LCLS-II

NCEM

The proliferation of accelerators and new technologies

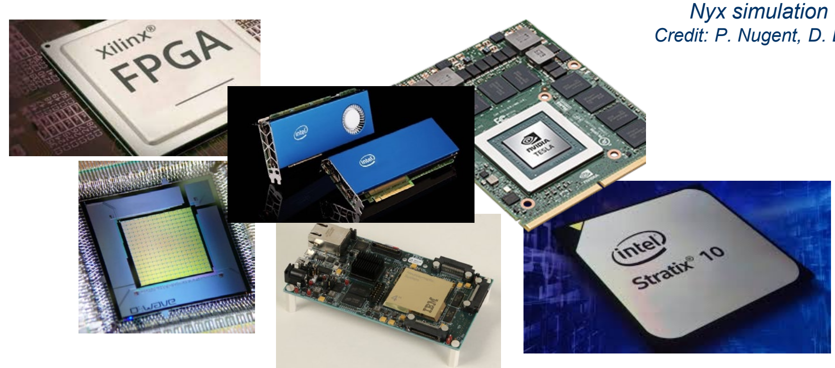

Use of advanced data analytics and AI in simulations as well as for integration of multimodal data sets

GANs used to construct surrogate models for creating weak lensing convergence maps

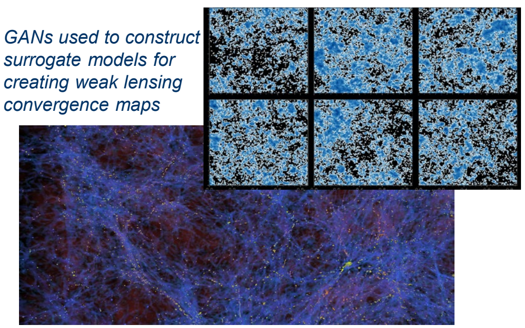

Nyx simulation of Lyman alpha forest
Credit: P. Nugent, D. Bard

From Katie Antypas et al.

BERKELEY LAB

NERSC has addressed the paradigm shift in the way we design, configure and operate HPC systems

Updated! NERSC-10 CD-0 achieved Sept. 2021

Users require an ***integrated ecosystem*** that supports new paradigms for ***data analysis*** *with* ***real-time interactive feedback*** *between experiments and simulations. Users need the ability to* ***search, analyze, reuse,*** *and* ***combine data*** *from* ***different sources*** *into* ***large scale simulations and AI models.***

NERSC-10 Mission Need Statement: The NERSC-10 system will accelerate end-to-end DOE SC workflows and enable new modes of scientific discovery through the integration of experiment, data analysis, and simulation.

NERSC-10 Architecture: Designed to support complex simulation and data analysis workflows at high performance

NERSC-10 will provide on-demand, dynamically composable, and resilient workflows across heterogeneous elements within NERSC and extending to the edge of experimental facilities and other user endpoints

Complexity and heterogeneity managed using complementary technologies

- **Programmable infrastructure**: avoid downfalls of one-size-fits-all, monolithic architecture
- **AI and automation**: sensible selection of default behaviours to reduce complexity for users

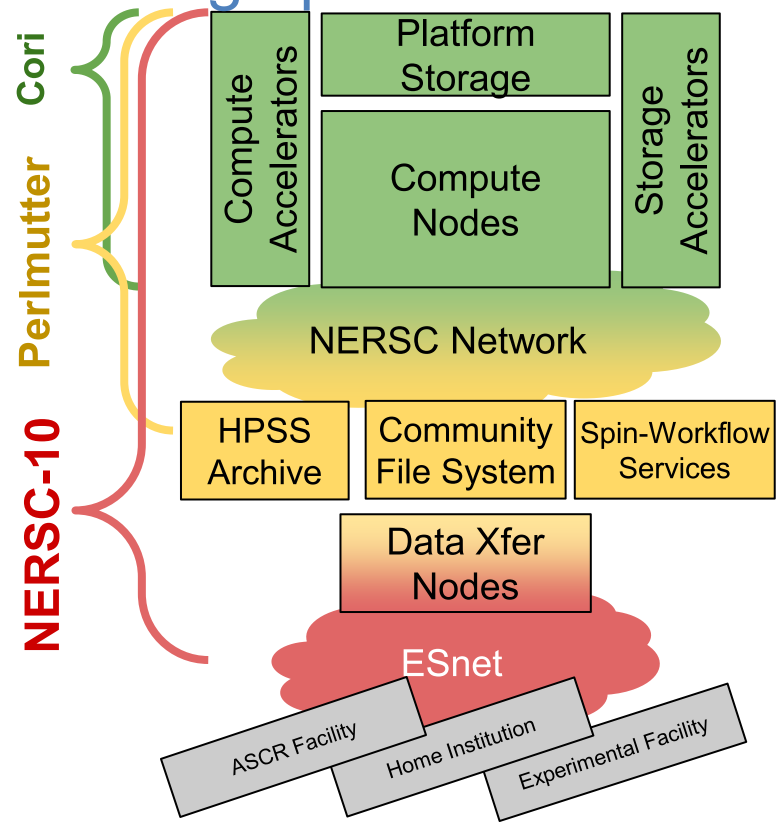

The new HPC environment needs …

- a catchy name (think back how “grid” stimulated research)
- a different set of benchmarks (focus on workflow and productivity)
- a different vendor relationship (not a single integrator)
- long term development cycle (not fixed milestones and a five year procurement effort)

And what about the "cloud"?

- National Labs have a mission to maintain a national capability
- We cannot allow the national R&D capability for HPC and computational science being taken over by commercial interest
- And it will be difficult to accomplish the tight integration with the cloud vendors

EXTRA SLIDES