Acceleration

Jay Kerley

Group VP & CIO, Applied Materials, Inc.

New Street Accelerated Computing Conference

OUR VISION

Our innovations **make possible** the technology shaping the future

OUR MISSION

To lead the world with **materials engineering** solutions that enable customers to transform possibilities into reality

Semiconductor Systems

Applied Global Services

Display and Adjacent Markets

The Scientific Method

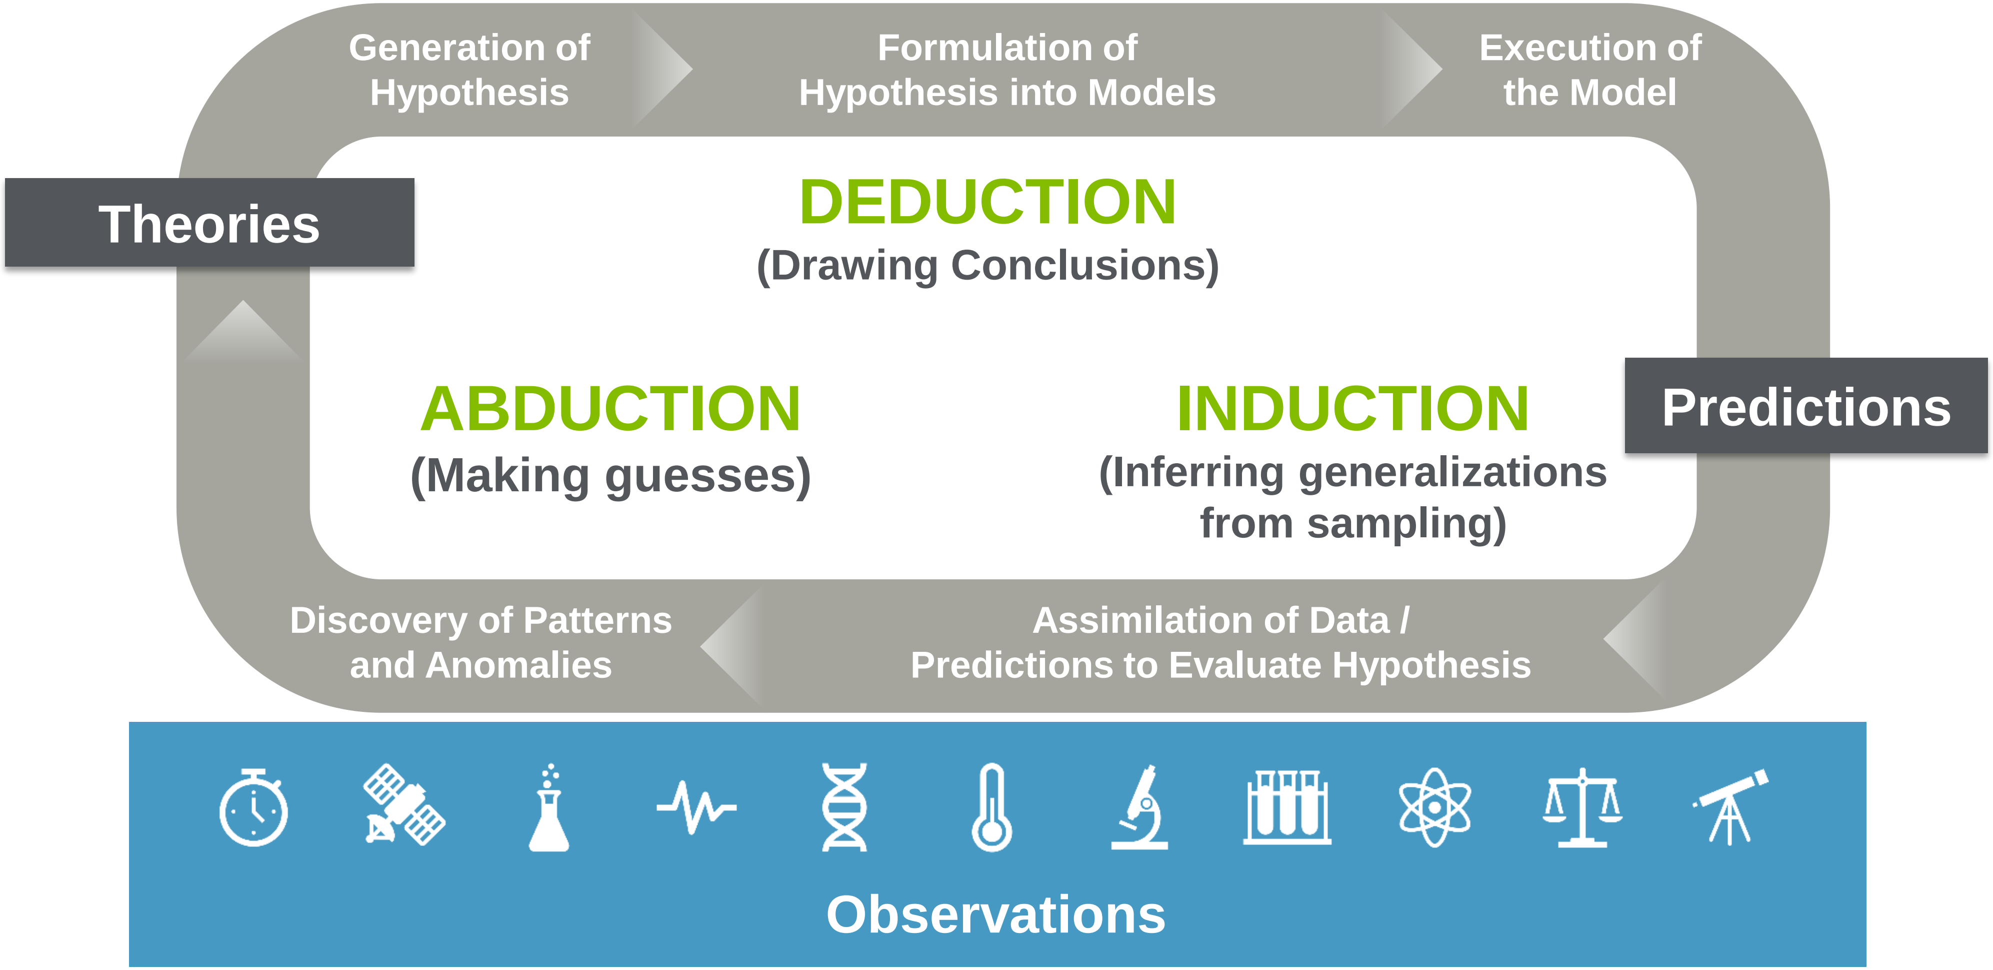

Source: Big Data and Extreme-Scale Computing Pathways to Convergence Report,

The Digital Scientific Method

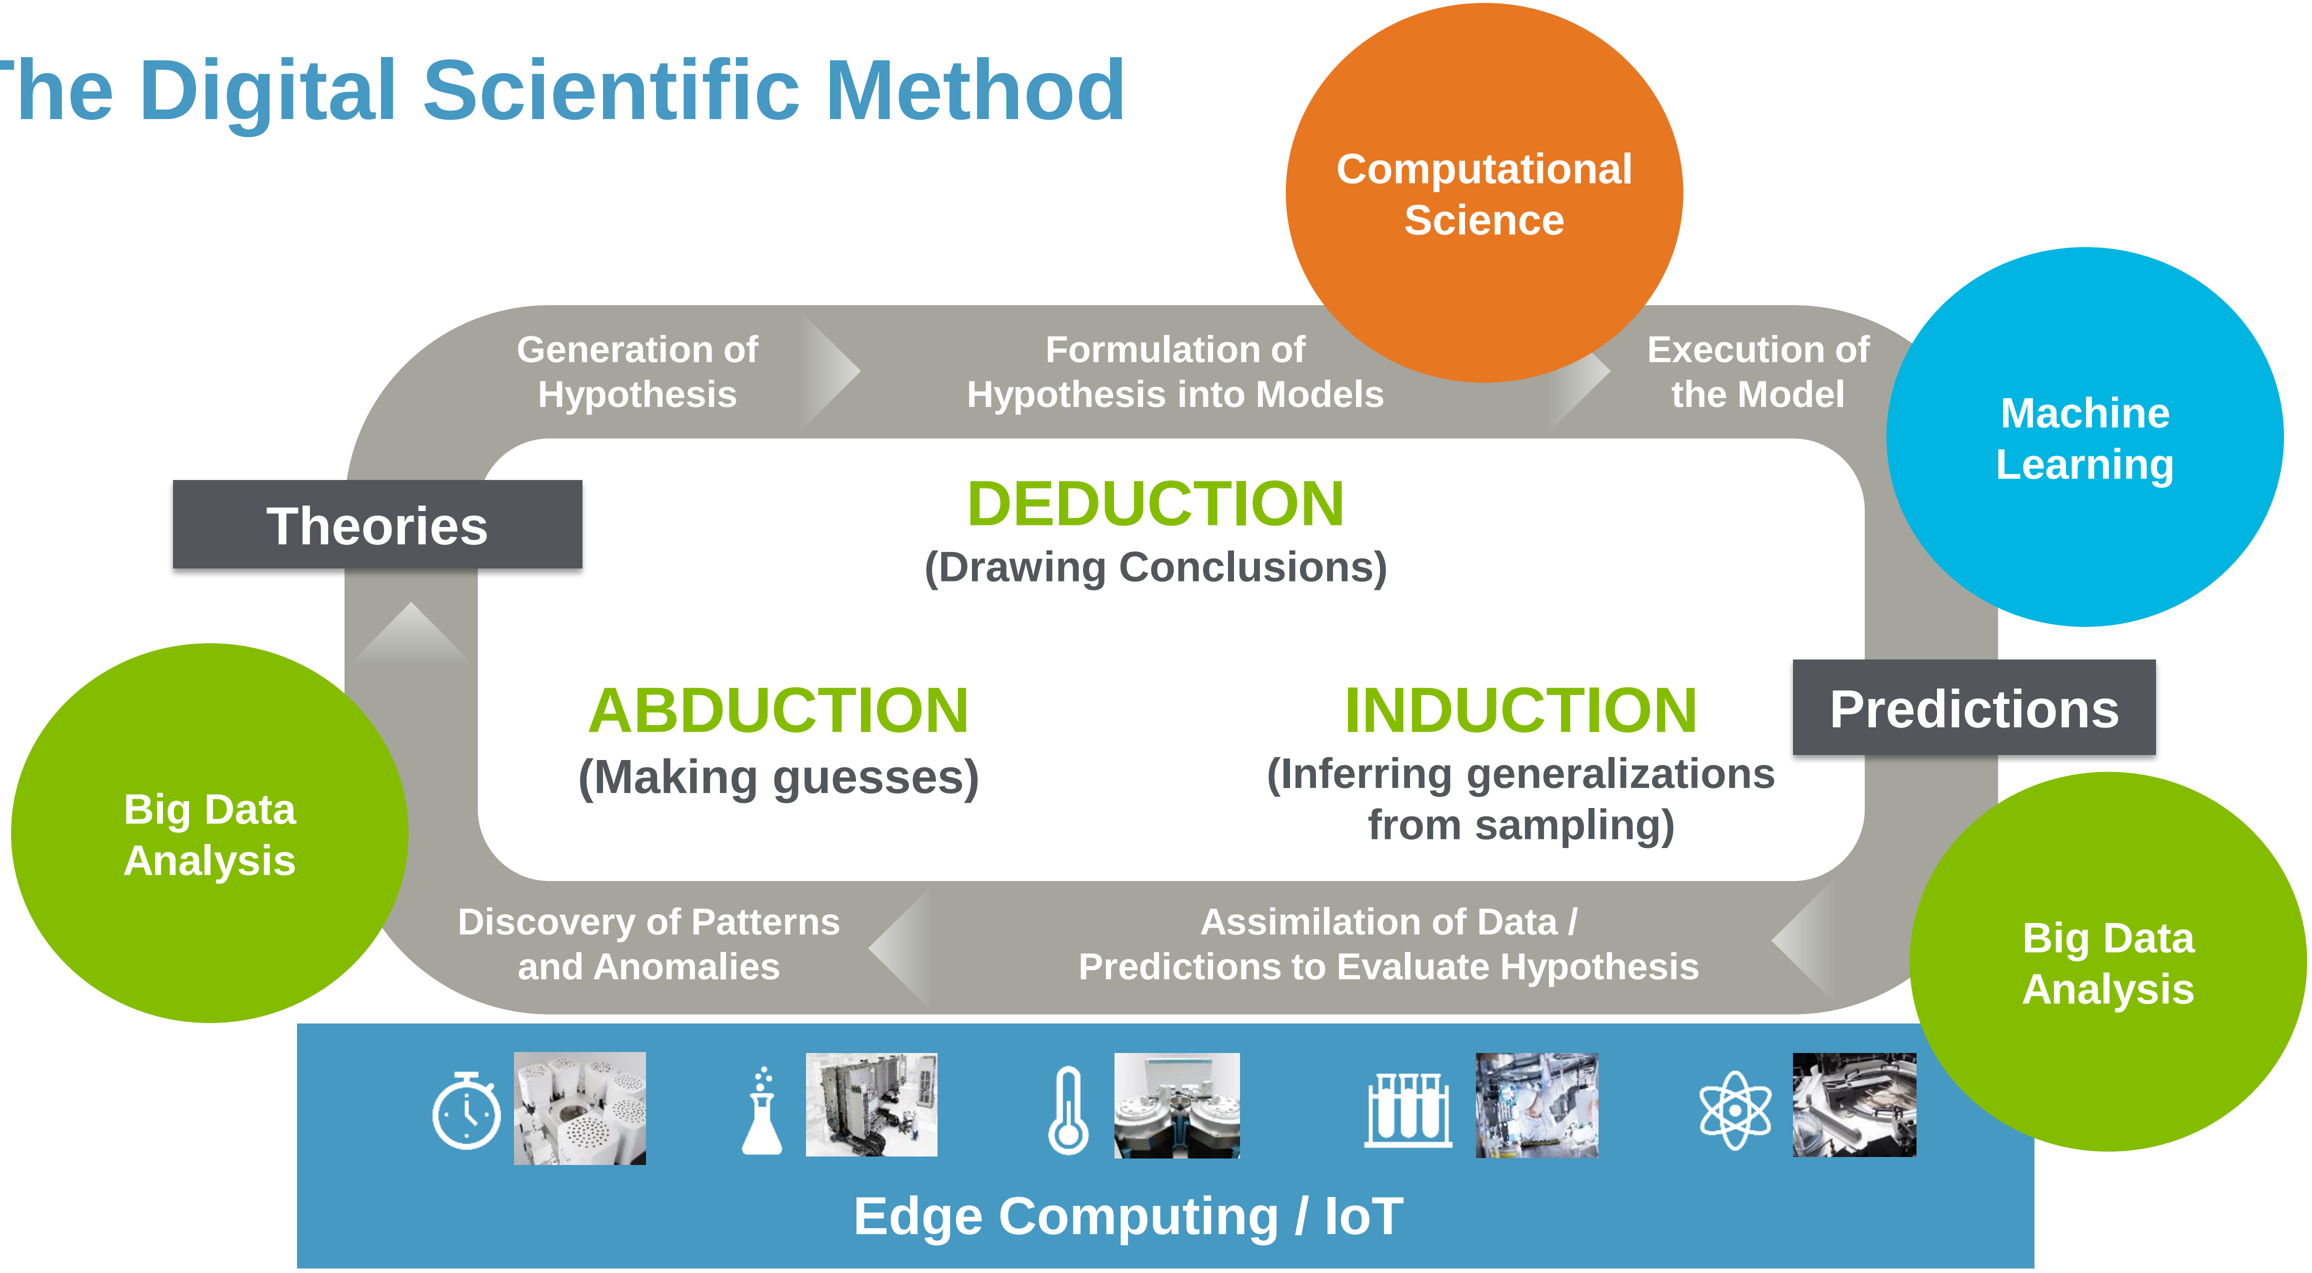

Source: Big Data and Extreme-Scale Computing Pathways to Convergence Report,

APPLIED MATERIALS®

It’s All About – Acceleration

Computational Science (aka HPC)	Big Data Analytics	Machine / Deep Learning
Simulation / Physics	Expansive Joins & Enrichment	Supervised
3D Visualization	High Speed Ingestion & Extraction	Unsupervised
Materials Modeling	Expansive Modeling	Reinforcement

Computational Science at Applied Materials

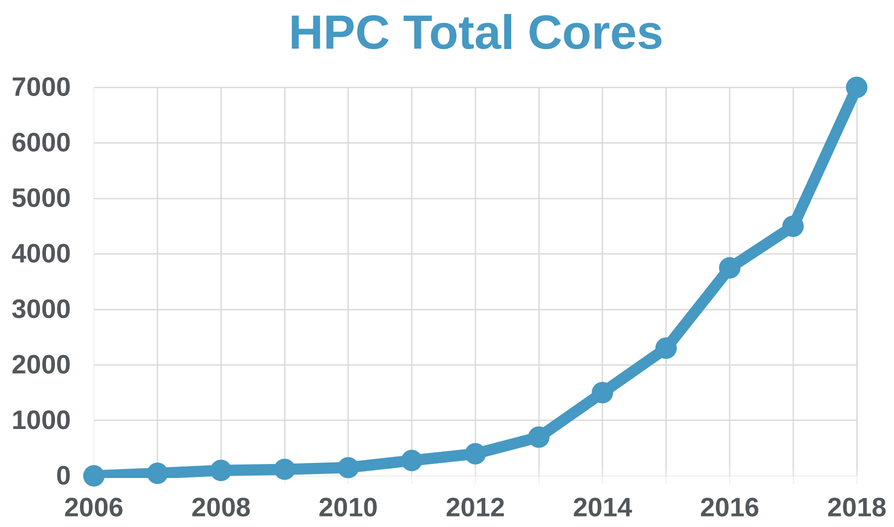

4496
Physical Cores

25TB
Memory

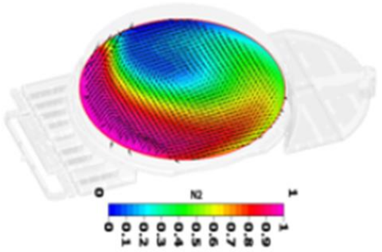

Flow/Chemistry

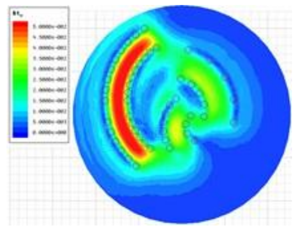

Magnetic Field

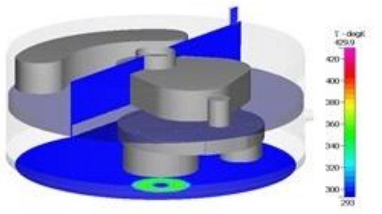

Thermal

Computational Chemistry

Surface engineering

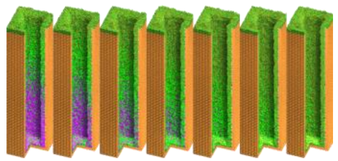

Film profile growth

Multi-scale, multi-physics modeling enabling hardware and process innovations at AMAT

SPEED Learning Rate and Time-to-Market

3D Visualization is Critical to Learning

$$SiHCl_3 \xrightarrow{k_{ad}} {}^{*}SiCl_2 + HCl\uparrow$$

$$^{*}SiCl_2 + H_2 \xrightarrow{k_r} Si(s) + 2HCl\uparrow$$

$$R_1 = k_{ad}[SiHCl_3]$$

$$R_2 = k_r[SiCl_2][H_2]$$

$$k_i = A_i T^{\beta_i} e^{-E_{a,i}/RT}$$

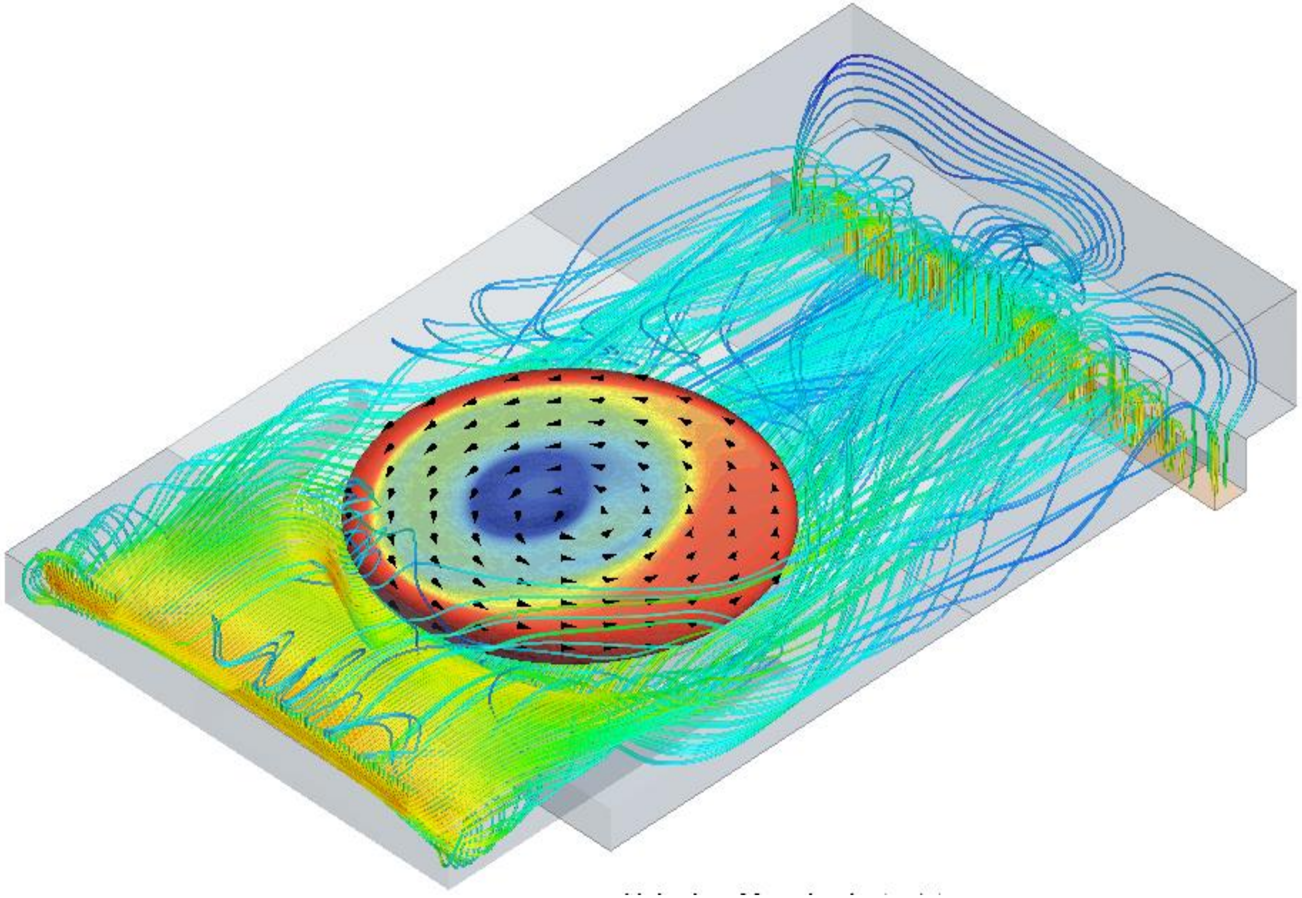

Translation into Visual Models

Source: Siemens AG 2017

Global Design Backbone

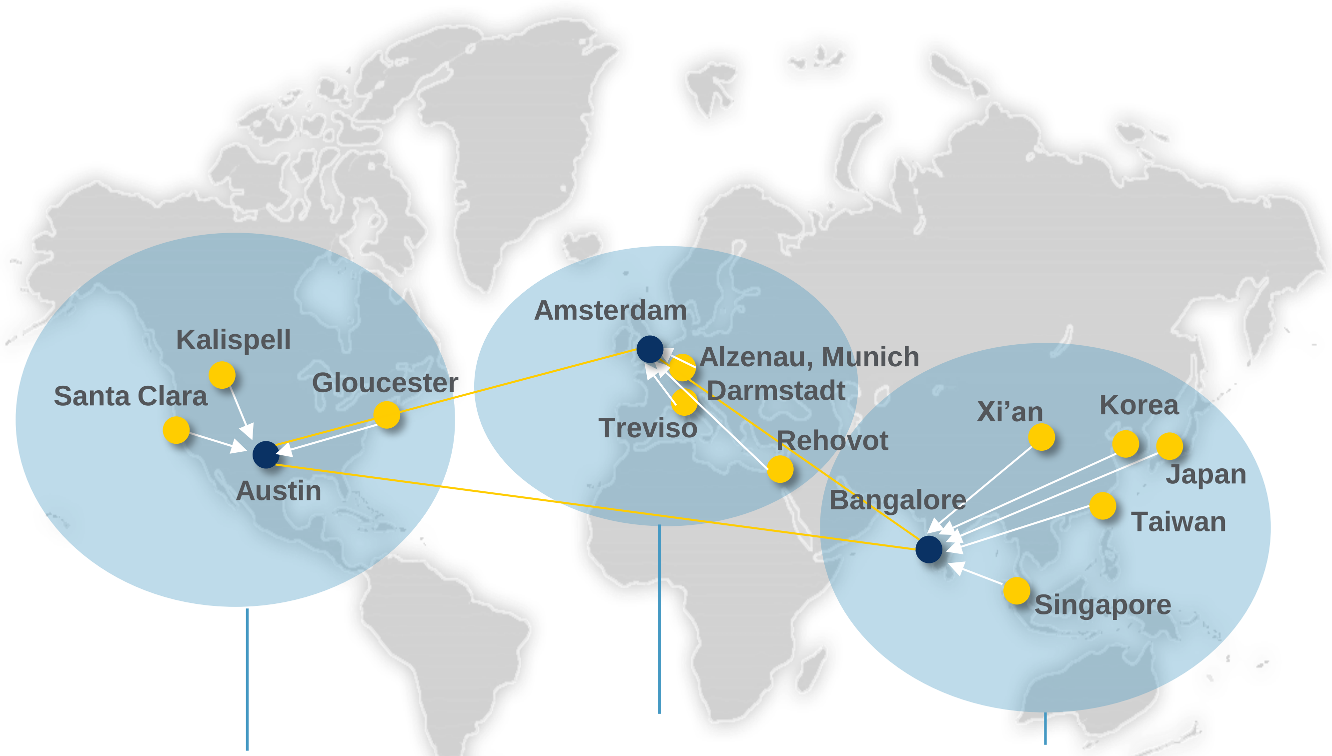

North America
~850 users

Europe
~150 users

Asia
~1000 users

1.6M CUDA Cores

> 3300 sessions

35,500 x86 HT Cores

> 877 Physical Blades

100TB of DRAM

GPU Cores

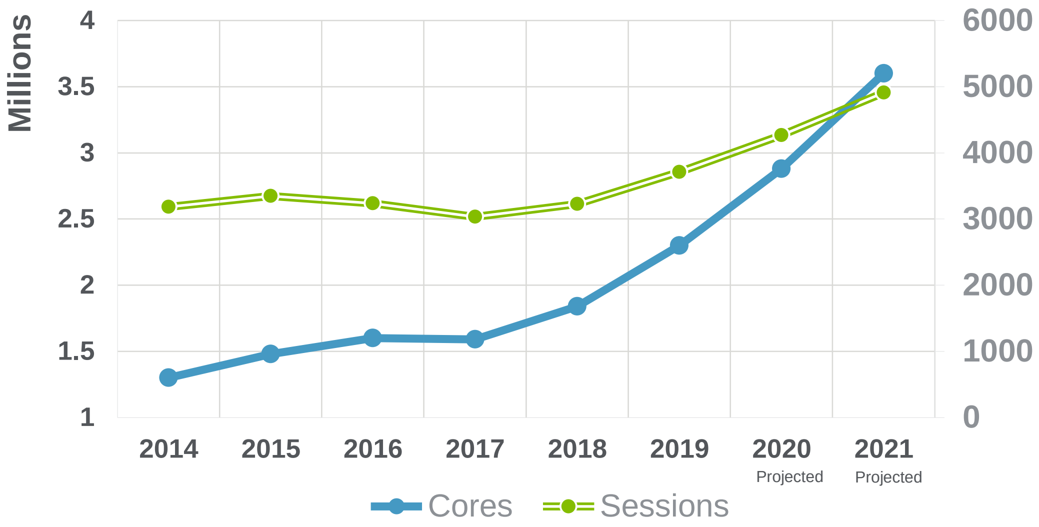

Implemented in 2008 – Award Winning GPU Cloud

Big Data in Semiconductor Manufacturing

Intelligent Control Requires Sensor Knowledge & Analytics

300mm
Single tool

50KB
data per second per tool

>500
tools per fab*

~1PB
annual data per fab

1,000,000,000,000,000

* 30K wspm fab

No. of Process Steps Grows 3x with Nodes

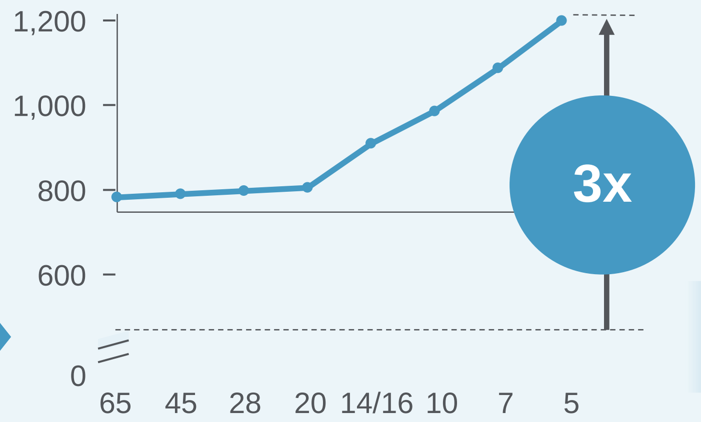

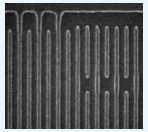

2D Device Architectures

3D Device Architectures

Quadrillions of Data Points by 2018

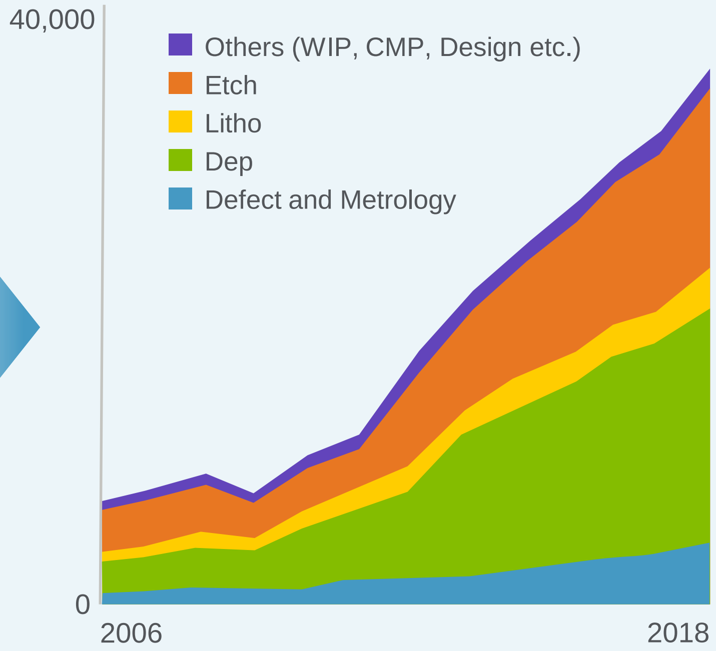

Source: VLSI Research & Applied Materials

How to Make Sense of the Big Data?

Applied Big Data Infrastructure

Data Sources

- Equipment Process Data
- Metrology
- Service History
- Manufacturing / Quality / Supply Chain
- Experiments
- Operational Data

Big Data Infrastructure

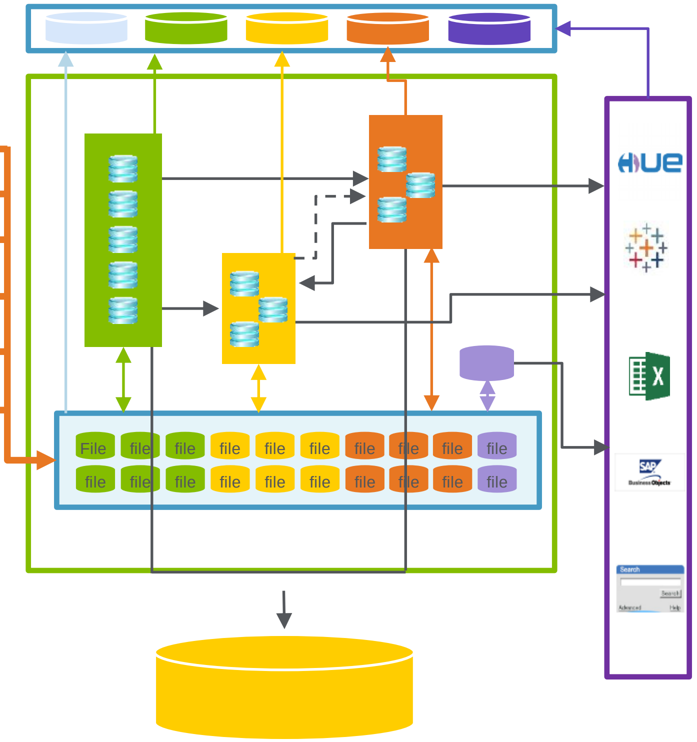

Technical Data Warehouse (TDW)

4.5PB Storage

66TB Memory

5000 Physical Cores

SAP HANA **7.5TB** Memory

Big Data Number of Cores

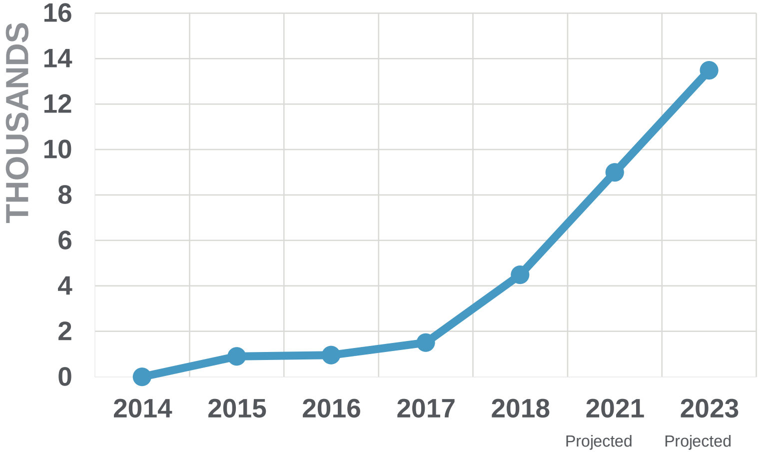

Machine Learning at Applied

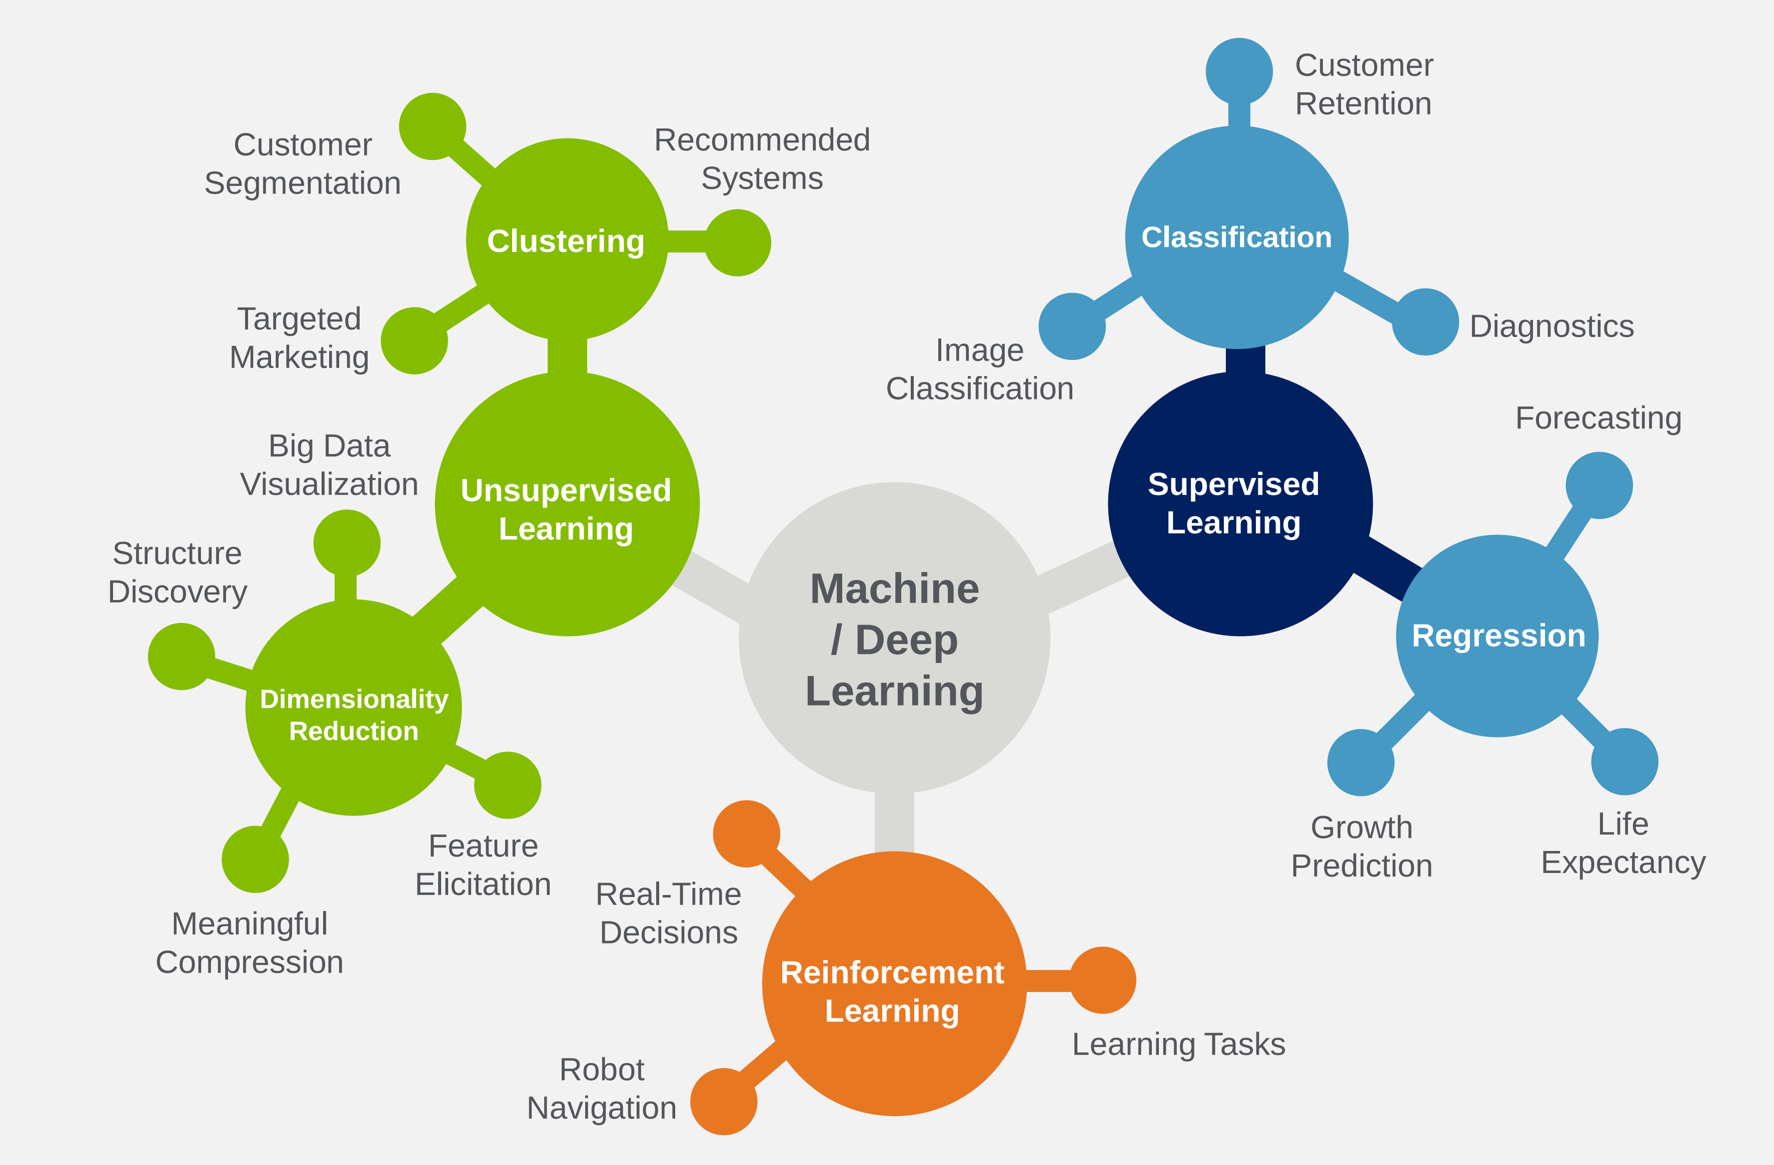

143K CUDA Cores

17.9K Tensor Cores

896GB Memory

CUDA Cores

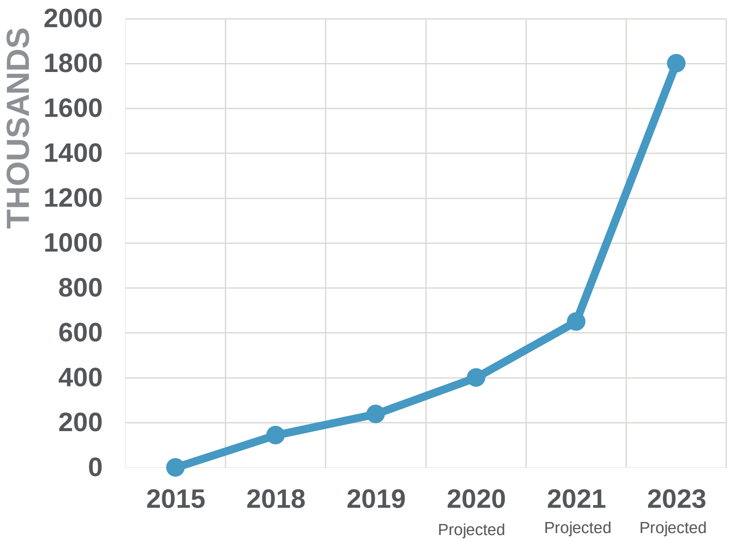

Testing, Trials and Options - So Many Options

Complexity in the data center is on the rise

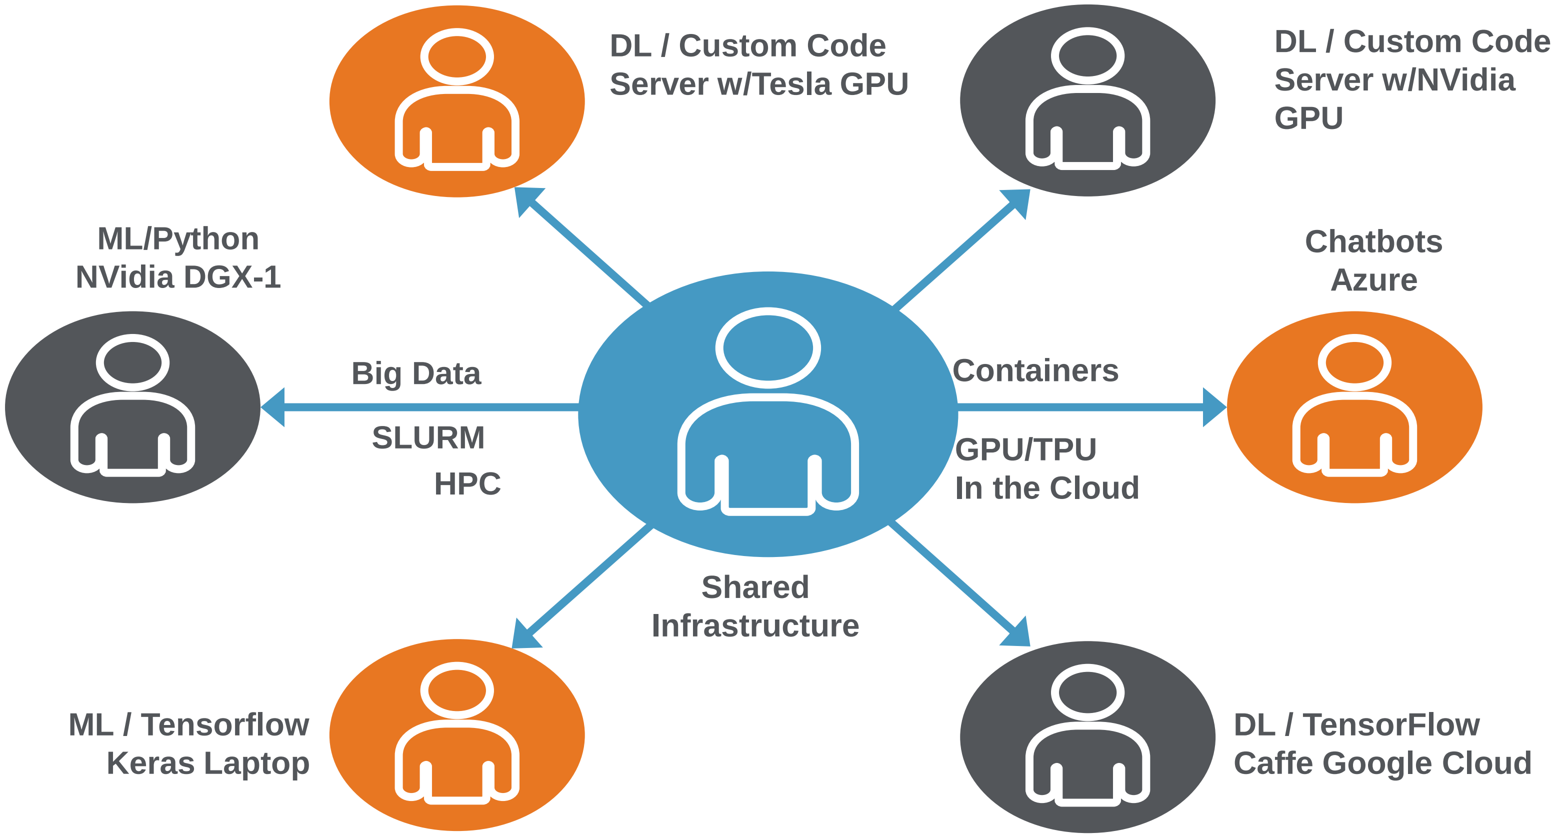

The Use Cases are Many – Still Early

APPLIED MATERIALS

BREADTH OF CAPABILITY AND PRODUCTS

Materials Creation	Materials Removal	Materials Modification	Inspection and Process Control
CVD, PVD, Epi, ALD, Plating	Etch, SRP, CMP, Cleans	Implant, Thermal	Optical, e-Beam

Connectivity
Broad Materials Solutions

\+

Acceleration
Computation Process Control

=

Valuable
Customer Results

APPLIED MATERIALS®

A System Level Approach Is Needed – It's Not Easy…

	Big Data Ecosystem	Comp. Science Ecosystem	Machine Learning Ecosystem
Application Level	Visualization, Mahout, R, Web Services	Applications and Opensource Codes	Speech, Image, Language
Middleware & Management	ZooKeeper (Orchestration); HIVE; PIG; Scoop; Flume; Map Reduce; Storm; Hbase BigTable; HDFS (Hadoop File Sys); HANA	C, C++, and other Development Env.; Domain Specific Libraries; Accelerator Tools; Numerical Libraries; Performance & Debug; Parallel FS; Batch Processor; Monitoring	TAP; Nervana; Spark; Caffe; Torch; Accelerator Tools; Numerical Libraries; Python; Parallel FS
System Software	Virtual Machines & Containers; LINUX OS Variant	Containers; GPU Virtual Machines; LINUX OS Variant	Accelerator Stack & Drivers; Open Fabric; LINUX OS Variant
Cluster Hardware	Ethernet Switching; Local Storage Node; X86 Racks & Memory Appliance	Infiniband Switching; SAN + Local Storage; GPU, X86, Memory Appliances	Advanced Interconnect & Infiniband; SAN + SSD Storage; GPU, FPGA,X86

Convergence is Needed - and so is Heterogenous Computing

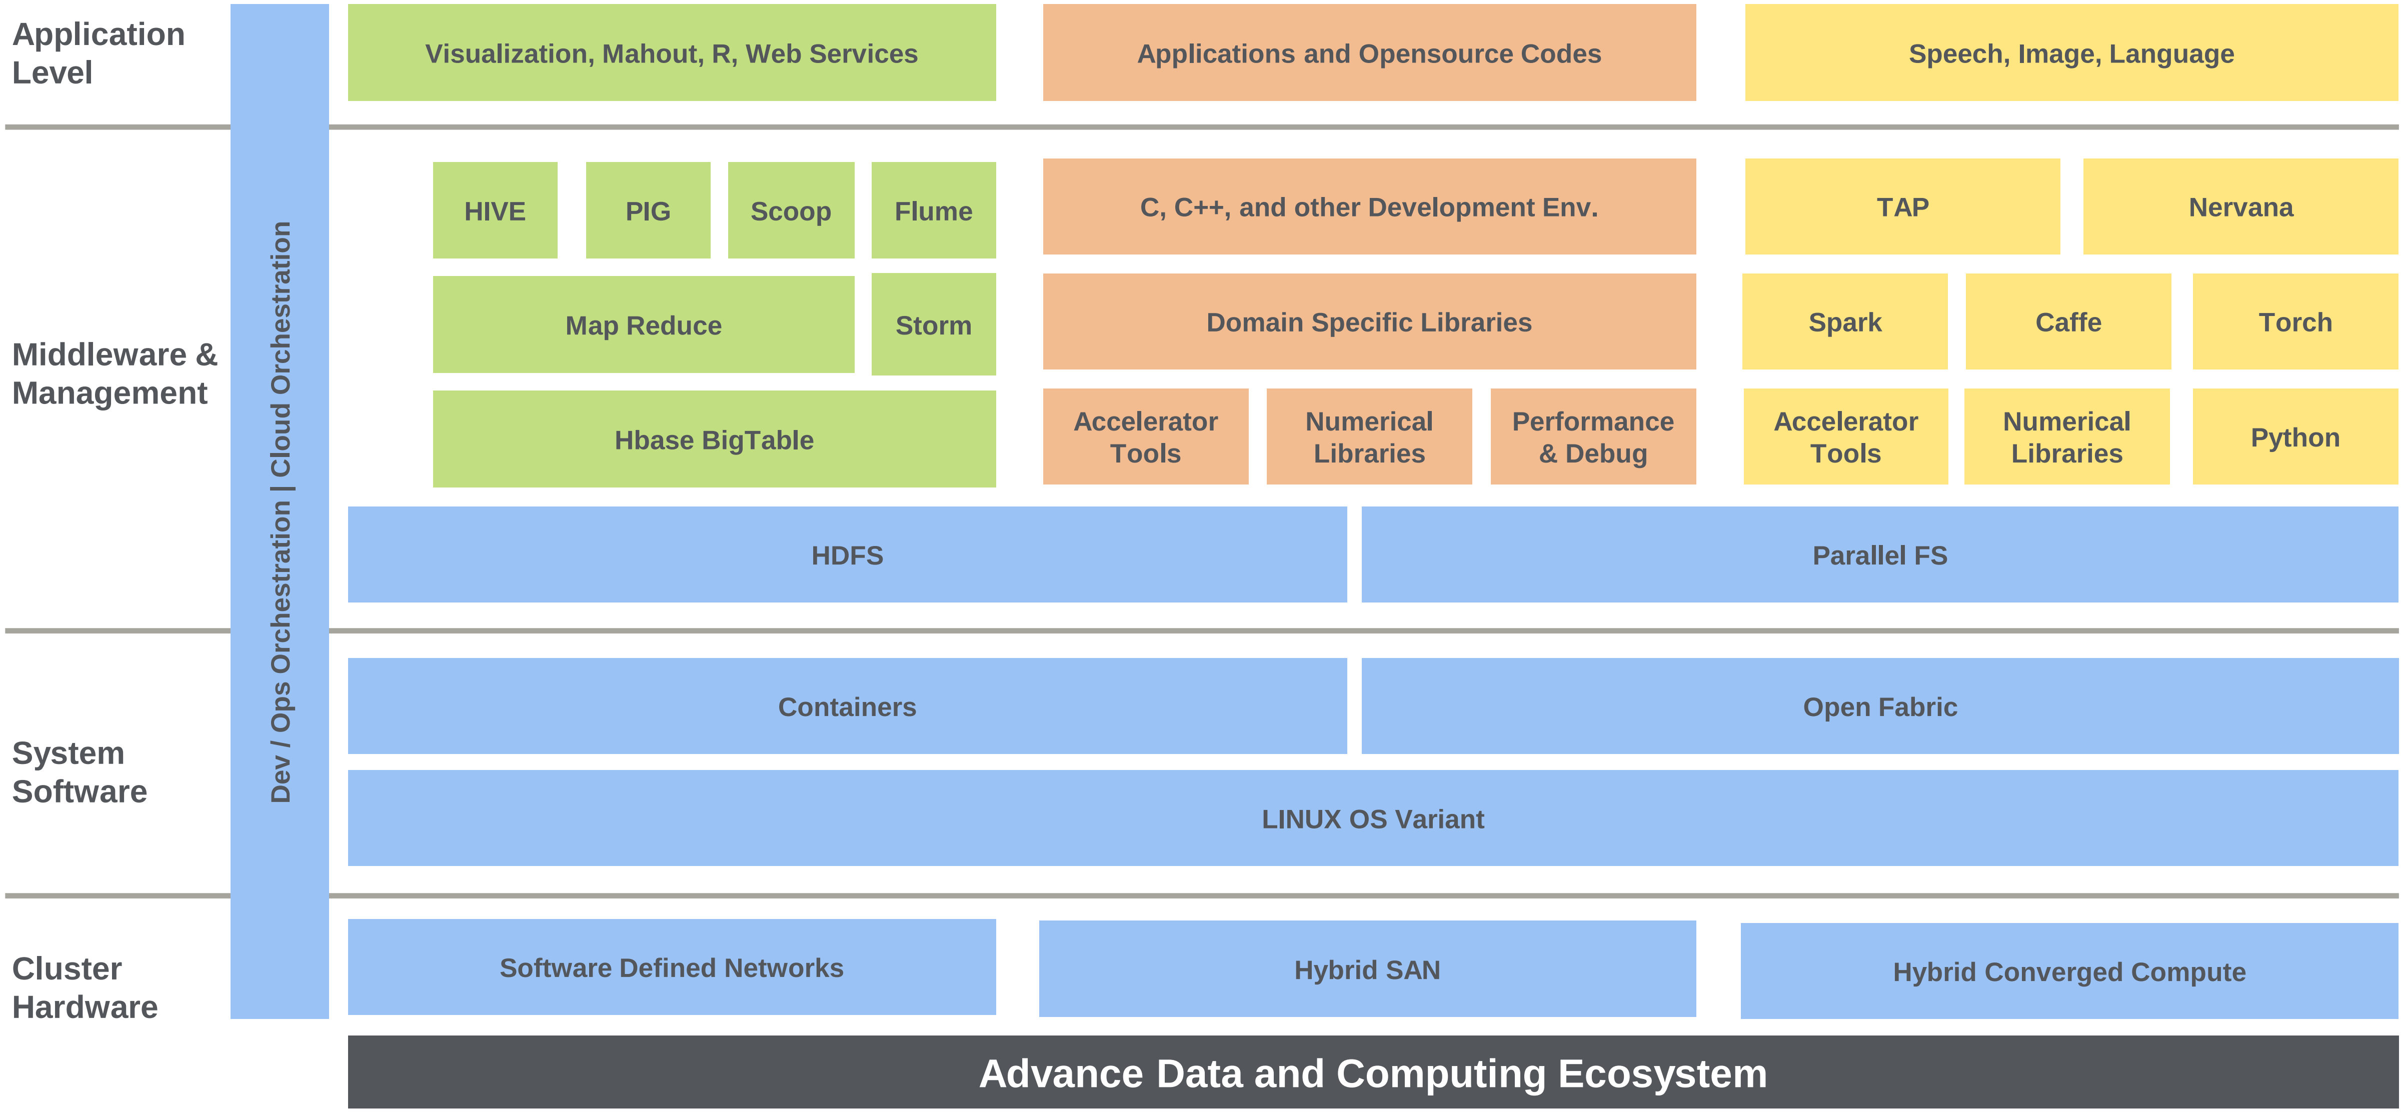

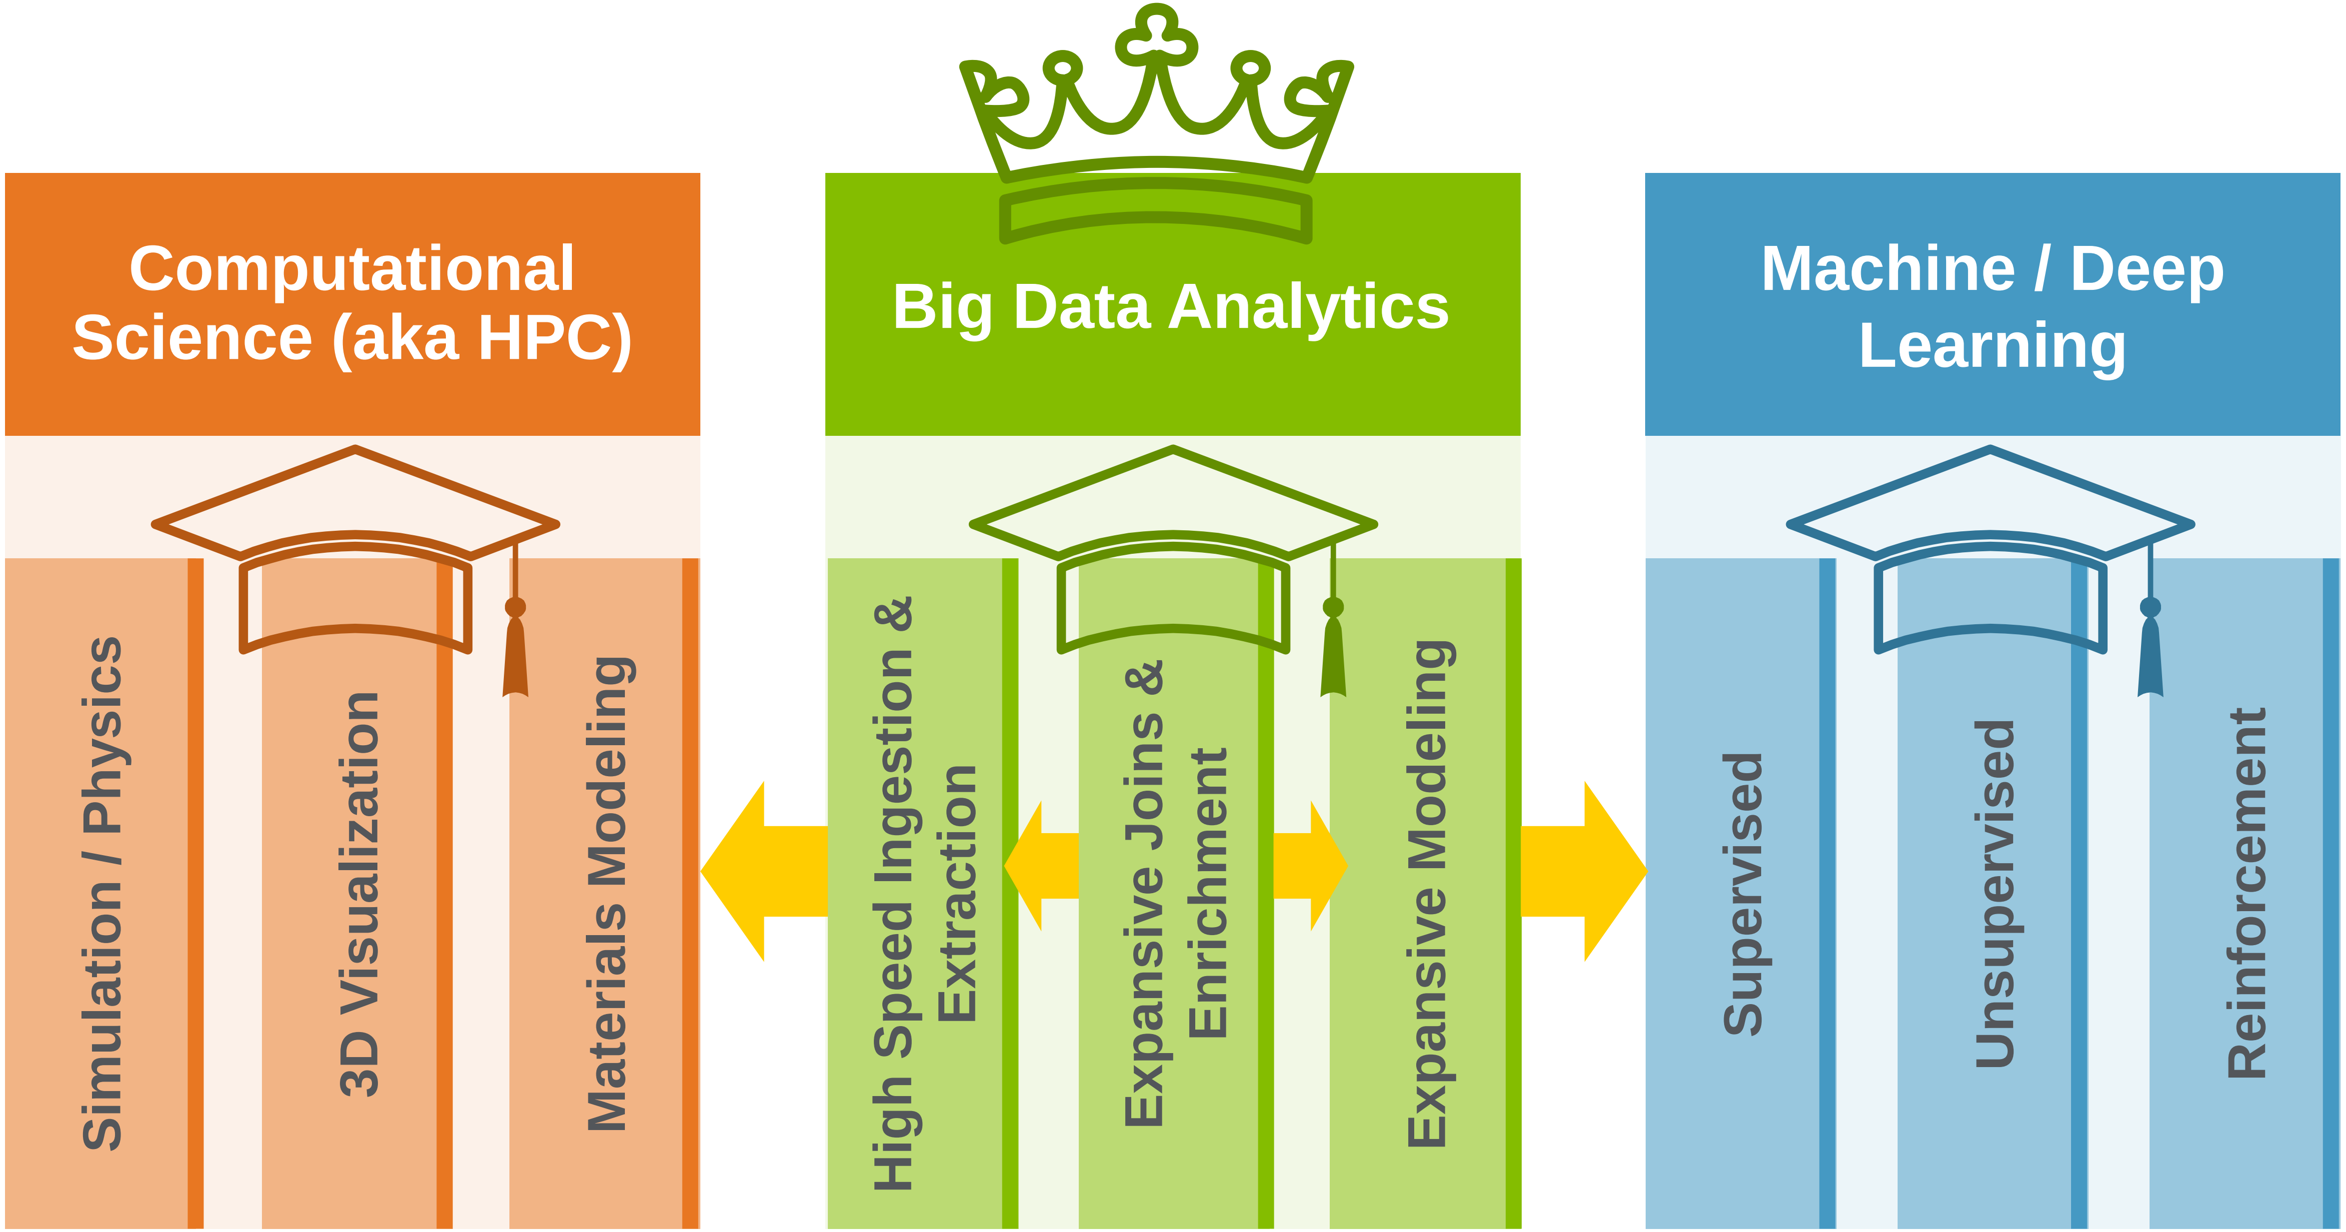
Computational Science (aka HPC)
Simulation / Physics
3D Visualization
Materials Modeling
Big Data Analytics
High Speed Ingestion & Extraction
Expansive Joins & Enrichment
Expansive Modeling
Machine / Deep Learning
Supervised
Unsupervised
Reinforcement